

Need to Know: Windows 8 Blue

Windows IT Pro

A PENTON PUBLICATION

MAY 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Translate Active Directory Object Names

Mailbox Auditing in
Exchange Server 2010

Windows Server 2012
Installation Options

What's New in Lync Server 2013



Mark Minasi's

Windows Power Tools

1&1 DYNAMIC CLOUD SERVER PRICE CONTROL



AS LOW AS

\$0.06
PER HOUR*



COMPLETE COST CONTROL

- **NEW: No setup fee!**
- **NEW: No base price!**
- **NEW: No minimum contract period!**
- **Full transparency** with accurate hourly billing
- **Parallels® Plesk Panel 11 included,** with unlimited domains



MAXIMUM FLEXIBILITY

- Configure the Processor Cores, RAM and Hard Disk Space
- Add up to 99 virtual machines



FULL ROOT ACCESS

- The complete functionality of a root server with dedicated resources



FAIL-SAFE SECURITY

- Redundant storage and mirrored processing units reliably protect your server



Parallels®
Plesk Panel

1&1



DOMAINS | E-MAIL | WEB HOSTING | eCOMMERCE | SERVERS


Call **1 (877) 461-2631** or buy online

1and1.com

* Other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Customer is billed monthly for minimum configuration (\$0.06/hour * 720 hours = \$43.20/month minimum). Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2013 1&1 Internet. All rights reserved.

CAN'T GET AWAY?

Get first-class education from your desk



Windows IT Pro offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to www.windowsitpro.com/events to see an up-to-date list of all online events.

WindowsIT Pro

First-class education from the top experts in the industry.

Visit www.windowsitpro.com/events for a knowledge upgrade today!

Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.

COVER STORY ▼

Translating Active Directory Object Names Between Formats

— Bill Stewart

29

The NameTranslate object is useful when you need to translate Active Directory object names between different formats, but it's awkward to use from PowerShell. Here's a PowerShell script that eliminates the awkwardness.

Features

39 Mailbox Auditing in Exchange Server 2010
Tony Redmond

53 Windows Server 2012 Installation Options
John Savill

65 What's New in Lync Server 2013
Byron O. Spurlock

Products

81 New & Improved

Interact

77 Ask the Experts

In Every Issue

85 Ctrl+Alt+Del

87 Advertiser Directory

87 Directory of Services

87 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns



6 [Need to Know](#)

Windows 8 Blue and How Microsoft Is Reinventing Itself

Paul Thurrott



11 [Windows Power Tools](#)

Where-Object Is the Filter of Filters

Mark Minasi



14 [Top 10](#)

PowerShell for Netsh

Michael Otey



17 [Enterprise Identity](#)

Peeking Into the Future of Identity

Sean Deuby



22 [What Would Microsoft Support Do?](#)

Understanding and Detecting Secure Channel Problems

Tim Springston

Editorial

Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
IT Community Manager: Rod Trent
Custom Group Editorial Director:
Dave Bernard
Exchange & Outlook, Cloud:
Brian Winstead
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialists: Karly Prickett &
Ashley Lawson

Advertising Sales

Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service
Senior Director, Marketing Analytics:
Tricia Syed

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, May 2013, Issue No. 225,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton Media, Inc. Copyright ©2013 Penton Media, Inc.
All rights reserved. No part of this publication may be
reproduced or distributed in any way without the written
consent of Penton Media, Inc.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2013, Penton Media, Inc., all rights
reserved. These programs may not be reproduced or
distributed in any form without permission in writing from
the publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton Media, Inc., under license from
owner. *Windows IT Pro* is an independent publication
not affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

Windows 8 Blue and How Microsoft Is Reinventing Itself



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



When Microsoft released its new generation platform releases in late 2012 and early 2013—[Windows 8](#) and Windows RT, of course, but also [Windows Server 2012](#), Windows Phone 8, the newly renamed Windows Services (formerly Windows Live), Office 365 and Office 2013, Visual Studio 2012, and more—it wasn't just iterating new versions of these products as it had in the past. Instead, Microsoft was undergoing a company-wide restructuring of how it delivers software solutions to its customers. Part of that restructuring involves changing how it services and updates these solutions going forward. This year, we're starting to see how that's going to work.

If you're familiar with Microsoft's traditional servicing schedule, especially for its business-oriented products, you know most Microsoft platforms have received monthly security updates, less frequent update roll-ups and other updates, and, once a year or so, a service pack. Microsoft had also briefly dabbled with something called a feature pack, which was essentially a way to deliver new features to existing products out of band with the major milestones, though more recently it also created "R2" (release 2) releases that accomplished a similar goal. And, of course, major new versions have appeared on some schedule as well, with one release roughly every three years or so.

This way of doing things made sense when most software was distributed on physical media and when the PCs used by office workers were overwhelmingly found within a corporate network where they could be easily managed and serviced. I don't have to go into detail about how the world has changed, but over the past decade, pervasive broadband, mobile computing, multi-touch devices, the consumerization of IT, the bring your own device (BYOD) to work movement,

and other technology trends have revolutionized how users expect to consume computing resources. Microsoft, often incorrectly perceived as a laggard, has aggressively moved to address these expectations.

The start of this was the addition of online services to the firm's stable of solutions. These include services that are unique to the online world—such as SkyDrive cloud storage—but also services-based versions of Microsoft's bread-and-butter on-premises offerings: Exchange Online, SharePoint Online, Lync Online, Office 365, and more. It includes new products that are like existing solutions but reimagined for the cloud: Windows Azure, for example, and Windows Intune.

Microsoft's strength is that it offers hybrid solutions in addition to its cloud-based and on-premises offerings. So whereas Google pretty much can only offer you Google Apps in the cloud, Microsoft can offer Exchange on-premises, Exchange Online (hosted by Microsoft or various partners), or a combination of the two. This flexibility has helped many companies plot a move to the cloud where and when it makes sense, on their own schedules.

But some of Microsoft's biggest products continue to be on-premises client solutions: Windows, of course, but also Office, each with its billion-plus installed base of active users. Surely there's no way to move such products to an online services model? Actually, there is.

Bringing Traditional Software into the Services Age

If you're familiar with Google's Chrome web browser, you might know that this product has iterated through more than 25 versions in its four and a half years of life. Compare this to Internet Explorer, Microsoft's web browser, which launched in 1995—18 years ago—and is only at version 10. In these two products you can see the difference between the online services way of doing things and the way in which Microsoft has traditionally updated its software. Chrome isn't an online service per se. It's traditional Windows software that downloads from the web and is installed locally on your PC, just like any other Windows application. But Google updates Chrome like

it's an online service. And users reap the benefits of a product that's automatically updated for them, quickly, with new features and bug and security fixes. It's basically a hybrid approach.

Windows 8 (and RT) and Office 2013 are two recent examples of traditional Microsoft software that's designed for this same servicing model. Both—not coincidentally—can be installed from a website and are in fact designed to be installed that way. Both install more quickly than their predecessors. And both are designed to be updated on a rolling basis. That's what we're starting to see in 2013.

This effort didn't develop overnight. In fact, Windows 8 and Office 2013—and other similarly designed products—are the results of years of change. One gets the notion that in the case of the Windows team, especially, change is coming a bit hard. But there is a plan.

First, most of the Microsoft products I've mentioned so far will be updated on a regular basis, like an online service, whether they're online services or not. (The schedules vary by product.) Monthly security patches will continue but service packs are out.

Those customers who receive Office 2013 via an Office 365 subscription—personal or work-based—will receive updates to the suite (as well as to the Office 365 hosted online services) on a quarterly basis. Yes, new features will still be delivered to Office 2013—including subsequent major releases, which one might still expect to occur on the old schedule—each quarter.

For Windows 8 and Windows RT, the plan is for the built-in apps to be updated on a rolling basis. The next release of Windows 8 and Windows RT—a sort of combination service pack and feature pack code-named Blue—will arrive in the second half of 2013, or at roughly GA + 1 (that is, one year after the general availability of the original RTM release of the OSs). You might think of Windows 8 Blue as being analogous to Windows XP SP2 or Windows Server 2008 R2, which Microsoft in the past might have charged customers for and named accordingly. But my sources tell me that this release will still be considered Windows 8 and that everyone will be moved forward.

The Windows 8/RT platform isn't the only one getting a Blue update in the coming months: Windows Server 2012, Windows Phone 8, and Windows Services are all getting their own updates as part of the Blue wave. And Office? Office 2013 is in fact getting an update, code-named Gemini, in late 2013 that will allegedly add Metro-style Office apps (Word, Excel, and PowerPoint) for Windows 8 and Windows RT, as well as Android and Apple iPad and iPhone versions of the suite.

Microsoft is expected to begin divulging information about Blue, and Gemini, and the next version of the Xbox, code-named Durango, which, get this, will be based on Windows 8 and will use the same developer APIs and frameworks. This could come soon—at TechEd 2013 in early June or at the recently announced Build 2013, which is happening in late June.

But you don't have to wait for that. I've got a few more details—in particular about Windows 8/RT Blue—courtesy of a leaked version of the product that I think provides a telling peek at Microsoft's plans for mobile clients.

Windows 8 Blue

As the second release of Windows 8, Blue moves the needle decidedly over to the Metro side of the dial. That is, while Windows 8 today is politely described as two OSs (Metro and the desktop) in one—and impolitely described as a Frankenstein's monster—Windows 8 Blue is all about reducing the reliance on the desktop.

The desktop's not going away in Blue, per se. My suspicion is that such a thing won't happen until Windows 9 at the earliest, could possibly happen in the Windows RT variant of the OS first, and will most likely be an optional change for a while. Still, one of the most confusing aspects of Windows 8 is that some settings are configured in the classic Control Panel interface while others are handled in the new PC Settings interface.

In Blue, there have been two major changes in this regard. First, PC Settings has been expanded, both with sub-screens and with more

configurable options; it's likely that by the time this OS update does ship, most users will be able to get everything they need exclusively on the Metro side. Second, many Metro-style settings are now available in the context of the interface. For example, when you view the Settings panel from the Start screen, you now actually see Start screen settings. In the initial Windows 8 version, you had to go spelunking into PC Settings to find that, and had to know about PC Settings to begin with.

What you don't see in the recently leaked build of Windows 8 Blue are any major changes to the desktop. There aren't any minor changes either: Blue appears to be about bolstering the Metro environment.

I'll rehash the argument that Windows 8 is already an excellent update for anyone using a traditional (non-touch) PC and that perhaps Blue doesn't need to make any headway in that particular area, whereas the Metro environment, as a 1.0 release, is obviously immature and in need of drastic change. But if you're nervous about this consumer-y vision of the future of computing, Windows 8 Blue is going to be a tough pill to swallow. In fact, I think it speaks pointedly to Microsoft's belief that multi-touch devices are indeed the future.

Or, perhaps more accurately, we might say that general-purpose computing isn't so traditional anymore. In addition to non-traditional PC devices, the next Xbox is taking the Windows 8 core to a living room device that will sport a new generation of natural UIs—motion and voice control—that could then easily be brought back around to PCs and even smartphones. It's a virtuous cycle just waiting to happen.

We'll find out more in the coming months. But between Blue, Gemini, and Microsoft's general plan to advance the state of the art in software deployment and servicing, it's pretty clear that the rusty old PC is about to get a big upgrade, or at least a ton of smaller upgrades. And it's happening, for a change, both inside and outside the box. ■

Where-Object Is the Filter of Filters

Winnow your PowerShell results to their essentials

You’ve already seen how using the pipeline (|) to string together a few PowerShell cmdlets can let you create a simple report, like the one I’ve discussed several times:

```
get-aduser -f * -pr lastlogondate | select
samaccountname,lastlogondate | sort lastlogondate
```

That’s a nice example of combining several separate PowerShell cmdlets (get the users, grab just their names and lastlogondates, and sort them by those lastlogondates), but it’s probably a little more than you really want. You don’t want to see *all* users and their last logon dates; you just want to see the ones who haven’t logged on since some unusually long-ago date.

For that, you would need a fourth PowerShell cmdlet, called *where-object*. Almost no one calls it *where-object*, though; most people use one of its shorter names, either *where* or simply the question mark (?). In past columns (such as in “[Doubling Up Active Directory PowerShell Cmdlets](#)” and later in “[3 PowerShell Account Tweaks](#)”) I’ve talked about using two-part cmdlets in PowerShell—the “filter” cmdlet (which finds just the user accounts that meet a certain criterion) and the “hammer” cmdlet (which performs an action against that discovered subset of AD users). You might also recall that I’ve discussed a handful of “filter” cmdlets: *get-aduser*, *search-adaccount*, and a few others.



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

Well, *where-object* is the filter of filters. Its syntax is a bit strange, however, so let's start with an example. To see all users who last logged on before January 1, 2013, you could type

```
get-aduser -f * | where {$_.lastlogondate -le "1 January 2013"}
```

First, *get-aduser* collects all the domain's user accounts, as you've seen before, and then it sends them down the pipeline, as you've also seen before. Then, *where-object* takes over. Its job is to examine every object emitted by the pipeline to see if that object meets some criterion. This criterion is written between the braces and yes, it's ugly, so let's pick it apart and make sense of it.

You'll remember that *-le* means "is less than or equal to," and the date on the right is obvious, but what is *\$_lastlogondate*? Why, it's the key to making a lot of one-liners work!

I've already said that *where-object*'s job is to test things arriving to it from the pipeline, and I think it's easy to see that a statement like *(something) -le "1 January 2013"* is a reasonable way for PowerShell to let you test *(something)* against a given date. But that *(something)* has to be whatever is in the pipeline at the moment, so you need some way to specify *whatever is in the pipeline at the moment*.

That *(something)* is called *\$_*. Yes, it looks strange, but here's the reasoning. I've written several times in other articles that PowerShell has variables, which let you store transient information in the computer's memory, and that you can immediately recognize that something is a variable because the first character in its name is a dollar sign (\$). Thus far, my examples of variables have been variables that I've created on the fly, but PowerShell—like most scripting environments—also includes a number of built-in variables. For example, there's *\$true* and *\$false*, built-in variables that store the values for *true* and *false*. Now, you'd think that PowerShell would store the current contents of the pipeline in a variable named *\$pipeline*, but it instead uses the name *\$_*. In this case, *\$_* contains an entire user account object.

But this article's sample criterion doesn't ask *if the user is less than January 1, 2013*, it asks if the `lastlogondate` property of that user is before that date. So, you need to be able to extract just the `lastlogondate` value, and that's why the criterion refers to `$_lastlogondate`. Adding a period and a property name pulls out just the part that you need. Here's what your one-liner would look like with *where-object* refining its output:

```
get-aduser -f * -pr lastlogondate | ? {$_lastlogondate -le
    "1 January 2012"} | select samaccountname,lastlogondate |
    sort lastlogondate
```

You can use *where-object* and `$_` to build all kinds of filters. For example, to see all users whose `samaccountname` starts with an *F*, you could type

```
get-aduser -f * | where {$_samaccountname -like "f*"}
```

So why did I spend so much time talking about the *-filter* option in *get-aduser*? Why not just make all one-liners look like

```
get-aduser -f * | where {$_<something> -like "<somepattern>"}
```

you might ask? From a technical point of view, there's nothing wrong with that, except for one thing: It's probably a massive waste of bandwidth and server time. The *get-aduser* cmdlet with a filter sends a command to a domain controller (DC) that allows a DC to return just a small subset of AD; *get-aduser -f ** piped into a *where-object* cmdlet tells the DC to deliver *all* the user accounts and then has the local CPU filter out the desired ones. *Where-object*, then, is a great general-purpose tool, but you should avoid it when there's a filter built in to your initial *get-whatever* cmdlet. ■

PowerShell for Netsh

Windows Server 2012 has vastly expanded PowerShell support



**Michael
Otey**

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

Email



Netsh has long been a mainstay among Microsoft Windows Server management commands, and it wasn't until the recent release of [Windows Server 2012](#)—with its vastly expanded PowerShell support—that PowerShell was able to do the same sort of tasks that you've always needed the venerable netsh to complete. In this Top 10, I'll show how to use PowerShell to perform 10 common tasks that previously required you to use netsh.

10 List network adapters—Perhaps the most basic network command is to list the network adapters. One nice bonus with the PowerShell command is that you can include the optional `-IncludeHidden` parameter to show hidden adapters that you can't see in the graphical interface. Here's the PowerShell cmdlet you need:

```
Get-NetAdapater
```

9 Disable and enable adapters—Another basic network adapter management function that PowerShell 3.0 allows is enabling and disabling the network adapter. Although I rarely do this on servers, it's pretty handy on laptops and tablets that have trouble connecting to different networks. The following commands will do the job for you:

```
Disable-NetAdapter -Name "Wireless Network Connection"
Enable-NetAdapter -Name "Wireless Network Connection"
Get-NetIPConfiguration -Detailed
```

The most basic Windows Server commands are to enable and disable Windows Firewall.

⑧ Show the system's TCP/IP information—The `Get-NetIPConfiguration` cmdlet shows the system's current TCP/IP configuration settings. You can see the system's current IPv4 and IPv6 addresses, the gateway address and its status, and DNS server addresses. Use the following command:

```
Get-NetIPConfiguration -Detailed
```

⑦ Rename a network adapter—You can rename network adapters by using PowerShell 3.0. However, you should be aware that you need to have administrative rights in order to execute this and other network configuration commands. Here's the command:

```
Rename-NetAdapter -Name "Ethernet" -NewName "Public"
```

⑥ Set a new static IP address—To set the IP address of the network adapter named Ethernet to 192.168.100.115 and the gateway address to 192.168.100.1, use the following type of command—it's especially useful when you're configuring Windows Server Core:

```
$netadapter = Get-NetAdapter -Name Ethernet
$netadapter | New-NetIPAddress -IPAddress 192.168.100.115
               -PrefixLength 24 -DefaultGateway 192.168.100.1
```

⑤ Set a new DNS server address—When you change the system's IP address type, you often have to change the DNS servers as well. The following example shows how the `Set-DNSClientServerAddress` cmdlet can be used to configure a network adapter with a DNS server address of 192.168.100.8:

```
$netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-DNSClientServerAddress -ServerAddresses
               192.168.100.8
```


④ Change the network adapter to use DHCP—PowerShell can be used to reconfigure the system to use a DHCP assigned address. This command sets the IP address of the network adapter named Ethernet to use DHCP addressing:

```
$netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-NetIPInterface -Dhcp Enabled
```

③ Set the Ethernet interface to use a DHCP assigned DNS address—When you switch to DHCP addressing, you typically want the DNS server address to be dynamically assigned. This command sets the Ethernet interface to use a DHCP assigned DNS address:

```
$netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-DnsClientServerAddress -InterfaceIndex 12
               -ResetServerAddresses
```

② Add a Windows Firewall rule—Just like netsh works with both the system's network configuration and the Windows Firewall configuration, the following PowerShell command shows how you can configure Windows Firewall to allow remote management:

```
Set-NetFirewallRule -DisplayGroup "Windows Firewall Remote
Management" -Enabled True
```

① Enable and disable Windows Firewall—The most basic Windows Server commands are to enable and disable Windows Firewall. The following commands show how PowerShell can do just that:

```
Set-NetFirewallProfile -Profile Domain,Public,Private -Enabled
True
Set-NetFirewallProfile -Profile Domain,Public,Private -Enabled
False
```

Peeking Into the Future of Identity

Simplicity for the user belies enormous complexity for the hybrid identity infrastructure

We're on the edge of an explosion in what we can do with cloud identity. As I mentioned in "[Reflections on RSA Conference 2013](#)," the most interesting session I attended at RSA was called "Emerging Conflicts In Identity Space," a panel discussion moderated by Todd Inskeep, senior associate at Booz Allen Hamilton. On the panel were some of the best-known thought leaders in cloud identity: Michael Barrett, [PayPal](#)'s chief information security officer; Andre Durand, [Ping Identity](#)'s founder and CEO; Chuck Mortimore, [Salesforce.com](#)'s vice president of product management, identity and security; and [Eric Sachs](#), Google's group product manager for identity. Across the session and all four panelists, several major themes emerged.

Passwords Aren't a Sustainable or Scalable Model

Although the identity community has been preaching for years that passwords aren't sustainable, it's become cliché to hear it. The problems with passwords have now become painfully apparent to everyone: Whether you're an identity expert or retired elementary schoolteacher, passwords are a woefully inadequate authentication solution today. Sure, we have to remember many passwords to log on to cloud service providers (CSPs) such as Facebook or Amazon. But the password problem is worse than that, because it's not a one-to-many (one desktop to many CSPs) relationship any more. It's a many-to-many relationship in which the user might be authenticating to these CSPs from a desktop, tablet, or mobile phone. And it gets worse! If you choose to use multifactor



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.



Email



Twitter



authentication (e.g., Gmail, Facebook), every app that accesses these CSPs requires its own, unique password—or at the very least requires that you input a PIN that has been sent to your cell phone. I can tell you from personal experience that using Gmail multifactor authentication across only Internet Explorer (IE) and Google Chrome browsers on my five separate clients (work notebook, desktop, Apple iPad, Microsoft Surface RT, Windows Phone 8) is a real pain. And good luck getting your PIN via text if you’re using Wi-Fi on an airplane!

Eric Sachs told a story about how his mother once described him as the “Internet Password Guy.” Over the past couple of years, though, she stopped introducing him this way. Mystified, he asked his mom why she stopped. She replied, “Do you really want to be known as that guy?” The anecdote received a good amount of laughter. But it was painful, knowing laughter, because the question resonates with us. When a relative or friend gets frustrated because of a forgotten password, I sometimes chime in and say, “I write about ways to fix that problem!”

Identity federation between identity providers (IdPs) and relying parties (RPs) is the accepted way to minimize the use of passwords,

and most large CSPs support it. Sachs had some interesting experiences to relate about what Google was seeing in federation adoption. First, CSPs—for example, Software as a Service (SaaS) vendors—were coming to him in increasing numbers, inquiring about how to convert their authentication infrastructure from a local user ID/password store to federating with an identity provider such as Google. Why such a rise in federation interest? No business wants to find itself on the front page of the *New York Times* as the latest company to have its user IDs and passwords published on the Internet.

However, Sachs said, it turns out that converting a CSP from an identity provider to a federation RP isn't easy, and there are relatively few tools available to make the job easier. ([Ping Identity's PingOne Application Provider Service](#) is the only tool I'm aware of, and it is itself a cloud service. Did I mention that identity is a confusing topic?) Sachs also said a surprisingly high percentage of enterprises were reverting back from federation. He attributed it to knee-jerk reactions from executives who couldn't access their SaaS apps due to failures in their on-premises federation service that might not have been designed with high availability in mind. To me, this is a strong selling point of IDaaS solutions, which are strongly designed for high availability and scalability.

Identity Needs to Move From the Perimeter to the Transaction

Mortimore focused on the need to change the way we think about the concept of the perimeter. "The perimeter only works when everyone is sitting at a desktop," he said. Traditional Identity and Access Management (IAM) systems are built around this notion, but they're being completely bypassed by mobile devices and cloud services. Not only are these devices working outside IAM systems; they aren't even the company's devices.

Mortimore stated that the new perimeter—and perhaps it should be called an *authentication point*—is identity plus context. Instead of a

clear, one-time challenge of user ID and password, it becomes a much more fluid process. Where are the users? What device are they on? Have they previously authenticated on this device? How many factors did they authentication with? What type of access is being requested in this transaction? This is the next step in authentication, variously called *transactional authentication* or *risk-based authentication*. It really is a radical departure from the old way of looking at authentication; Barrett said one reason he enjoys working at PayPal is that the company employs some degree of this method already.

The Identity World's Inflection Point

Durand described identity today as a collection of Lego building blocks. I've long used the idea of puzzle pieces to describe the complexity of this market, but I prefer the building blocks analogy, because unlike puzzle pieces, building blocks can be combined in many ways to construct many different things.

And so it is with the web of identity (see "[The Strands of Your Identity Web](#)"): All the moving parts—authentication, authorization, account provisioning, auditing and governance—can happen in many places, in many ways, with different protocols, at many different times. It all depends on how and where you're able to put the blocks together to support the use case.

Durand also used biology to describe the identity environment. Today, he said, it's mostly like a collection of single-cell organisms, in one-to-one relationships with each other. (A popular example is that of one enterprise IdP to a single RP/SaaS vendor.) You might have a number of these relationships, but they're all one-to-one.

Identity in the near future, however, will be more like a multi-cell organism in which the relationships won't be just one-to-one but many-to-many. It will include transactional authentication, biometrics, and a degree of portability that we just don't have today. An example might be a user authenticating to an enterprise with his or her Google or Facebook credentials, but also presenting other identity

attributes to that enterprise that are associated with the user's bank (e.g., a bank account number for the payroll department to use).

We're close to an inflection point, the knee of a curve at which all this identity flexibility will be available to secure a variety of use cases that the panel hadn't even thought of a few years ago. Almost all the building blocks—identity standards—are in place for this to happen, if they're supported and implemented. (Barrett emphasized the value of [FIDO Alliance](#), a narrowly targeted open-standards consortium on interoperable strong authentication that's also easy to use.) But this dramatically more complex identity infrastructure must be transparent to users, especially consumers, because they'll always opt for the easiest method regardless of its security. Barrett offered a great security design principle: "Consumers go for convenience over all else, and they simply expect security to come with it."

3 Steps to Identity's Future

Durand put forth a clear three-step vision of identity's future. The first is where we are today: extending your corporate identity to the cloud via on-premises federation. The second step is gaining momentum: using an IDaaS provider to do the heavy lifting of this identity extension by simplifying the trust connection between the enterprise and the rest of the world, handling the web of federated trusts to CSPs, providing high availability, and (in the short term) providing password vaulting to create single sign-on (SSO) to CSPs that still require direct accounts. The third step, which might not be realized for years, is portable identity. In this phase, "you have a cloud identity provider that is affiliated with you the individual," Durand said. "Your affiliation with corporations is temporary and is viewed as credentials being added to your base identity. When you have these corporate credentials, you can get into the corporation. And they're taken away when you terminate."

And we'll have eliminated all but a few passwords. Can I get an "Amen!" from everyone? ■

Understanding and Detecting Secure Channel Problems



Tim Springston

is a senior support escalation engineer on the Commercial Technical Support team at Microsoft, where he is the lead for security and authentication. Check out his [Active Directory Blog](#).

Email



Blog



The Microsoft domain infrastructure design has some complicated aspects. Active Directory (AD), for example, relies on a commonly defined and working schema for objects and attributes in the database, demands network connectivity to peer domain controllers (DCs) to ensure timeliness of item updates, and needs DNS configuration to be correct, as well as other environment dependencies.

Every computer that's joined to a domain—whether it's a workstation client, a server, or a DC—requires connectivity to DCs in order to fulfill some of the service requirements that AD domains require. For workstations and servers, that connectivity is to the DCs in the domain that they're a member of, as well as trusting domains' DCs. DCs in one domain need connectivity to DCs in other trusting and trusted domains. The name describing the cached values in that inter-domain connectivity is the “domain secure channel.” To be clear, there are two kinds of secure channels: secure channels from a domain member to a DC in its domain, and trust secure channels between a trusted and trusting DC.

Why Do Secure Channels Matter?

Why does someone in support care about the health of the secure channel? The reason is that all domain-related services rely on the secure channel to a greater or lesser extent. Can't get Group Policy? Check the secure channel. Can't access a network resource? Check the secure channel. Can't log on to the domain? Check the secure channel. Of course, there are other things that can cause these same problems, but few are more difficult to diagnose—or more common—than a problematic secure channel.

What do these computers do with these secure channels? The flip-pant answer is, “Anything domain-related, of course!” All domain-related services rely on the ability to locate a DC to send the service request to. This is true for a domain member (e.g., a workstation, a member server) as well as a DC. The availability of a responsive DC is really all that a secure channel is. If a server can’t be contacted to send the requests to, the services fail.

For example, if a user connects to a SharePoint site that’s configured to use Kerberos, he or she will need to request a Kerberos ticket to pass to that SharePoint server for authorization. The user’s computer looks to the cached secure channel information for that domain (a cache that’s maintained by the NetLogon service) for which DC to send the Kerberos ticket request to. If that DC is unresponsive for whatever reason, the ticket request won’t occur, and the SharePoint connection won’t be authenticated using Kerberos. Depending on the SharePoint configuration, this might result in a lack of access to that site—all because of a secure channel-related problem.

Let’s walk through a typical multi-domain scenario. Suppose User A from Domain A has logged on to Computer B in Domain B. A logon Group Policy for the user is processed, and a Domain A DC is queried via LDAP to determine which Group Policies are applicable to User A. How does Computer B, which is a member of Domain B, know where to send the network traffic in order to find out which Domain A policies to process? It’s able to do that because the network location information for that domain and a DC in it are kept up-to-date constantly. That information is kept up-to-date by the NetLogon service on every domain-joined Windows computer. The NetLogon service is constantly maintaining its list of available DCs and domains (when trusts are present). Figure 1 shows a NetLogon debug log snippet that shows some of that ongoing routine. Of course, you can view your own NetLogon debug log information on your computers by following the steps in the Microsoft article “[Enabling debug logging for the NetLogon service.](#)”

Figure 1: NetLogon Debug Log Snippet

```

03/26 14:51:04 [MISC] [19236] NetpDcGetName: americas.contoso.com. using cached
    information ( NtDcCacheEntry = 0x00000003887FC68D0 )
03/26 14:51:09 [MISC] [11528] DsGetDcName function called: client PID=1292,
    Dom:(null) Acct:(null) Flags: DS GC RET_DNS
03/26 14:51:09 [MISC] [11528] NetpDcInitializeContext: DSGETDC_VALID_FLAGS is
    c03ffff1
03/26 14:51:09 [MISC] [11528] NetpDcGetName: americas.contoso.com. using cached
    information ( NtDcCacheEntry = 0x00000003888470170 )
03/26 14:51:09 [MISC] [11528] DsGetDcName: results as follows: DCName:\\
    BYC-NA-DC-50.americas.contoso.com DCAAddress:\\<snipIP> DCAddrType:0x1
    DomainName:americas.contoso.com DnsForestName:contoso.com Flags:0xe00031fc
    DcSiteName:NYC-RHQ ClientSiteName:NYC-RHQ
03/26 14:51:09 [MISC] [11528] DsGetDcName function returns 0 (client PID=1292):
    Dom:(null) Acct:(null) Flags: DS GC RET_DNS

```

At a high level, the causes of a secure channel problem can be boiled down to network connectivity problems. If the connectivity problems are intermittent, then when the network is working, all services are working as well. If the connectivity problems persist, then they stand a likelihood of causing a *broken secure channel*. A broken secure channel just means that the shared secret between the computer and AD would be dissimilar, and as a consequence the computer would be untrusted. The net effect in that situation is that no one would be able to log on to the domain and gain access to domain resources.

On a client computer or member server, a broken secure channel is bad because it might affect that computer's authentication to network services and any other services it provides. On a DC, it could prevent AD replication and cause unexpected logon and access problems if left untreated.

Identifying a Secure Channel Problem

The best way to identify whether a computer is having a secure channel problem is to do something that ultimately calls the

I_NetLogonControl2 function. I_NetlogonControl2 is one of the functions used in the NetLogon service (which is present on any Windows computer in any OS version) to keep knowledge about what domains and which DCs are accessible.

An IT pro has three easy ways to call that function and get a quick “thumbs up” or “thumbs down” about connectivity to a specific domain and DC: NLTest.exe, PowerShell, and WMI.

NLTest.exe. NLTest.exe was shipped in the Windows 2000 and Windows Server 2003 Support Tools but is contained by default in most later Windows OSs. NLTest.exe’s *sc_verify* switch calls I_NetlogonControl2, and you simply need to supply the domain you’re concerned about.

```
C:\>nlttest /sc_verify:americas
Flags: b0 HAS_IP HAS_TIMESERV
Trusted DC Name \\DD3-AM-DC-03.americas.fabrikam.com
Trusted DC Connection Status Status = 0 0x0 NERR_Success
Trust Verification Status = 0 0x0 NERR_Success
The command completed successfully
```

In those cases where the secure channel problem can’t heal itself—when the computer’s shared secret is dissimilar to what AD has for that computer—the NLTest.exe *sc_reset* switch can repair the problem.

PowerShell. The Test-ComputerSecureChannel PowerShell cmdlet was added in PowerShell 2.0. This cmdlet also calls I_NetLogonControl2 but provides less detail in its test—it simply returns a Boolean response of True if the domain secure channel is healthy and a DC is reachable, and False if not.

```
PS C:\> Test-ComputerSecureChannel
True
```

Similar to NLTest.exe, Test-ComputerSecureChannel can be used to fix the problem as well by calling the Repair switch.

WMI. By using the `win32_ntdomain` class, Windows Management Instrumentation (WMI) can query all the domains that the computer knows of. WMI can be useful in situations in which you can't rely on PowerShell being installed on the computer you're testing. Note that the following sample (where `Win32_NTDomain` is called via PowerShell's `Get-WMIObject` cmdlet using its alias `GWMI`) shows only the local domain but would actually return every domain that the local domain has trusts with.

```
PS C:\> gwmi win32_ntdomain
ClientSiteName           : TX
DcSiteName               : TX
Description              : AMERICAS
DnsForestName            : americas.fabrikam.com
DomainControllerAddress  : \\<IPaddress>
DomainControllerName     : \\DD3-AM-DC-03
DomainName               : AMERICAS
Roles                   :
Status                  : OK
```

Note that the status of OK in that example corresponds with the `True` or `False` return from `Test-ComputerSecureChannel`.

Fixing a Secure Channel Problem

In the case of Microsoft Customer Service & Support's Commercial Technical Support, we send an optional data-collection package to customers who call us for support. Within that package, we use WMI's `Win32_NTDomain` class (called via PowerShell) rather than PowerShell's native `Test-ComputerSecureChannel` cmdlet, because we want to be sure that the test will run even on older OSs such as Windows XP and Windows 2003. What specifically do we do in our test? In Listing 1 and Listing 2, I present some script samples, using the same methods we use, that you can call via PowerShell on your own.

Listing 1: Obtain Domain Name Secure Channel Information for the Current Domain

```

PowerShell
Get-Date >> $OutputFileName
$ComputerName = Get-WmiObject -Class Win32_ComputerSystem
$OutputFileName = Join-Path $Pwd.Path ($ComputerName.Name + "_Secure
    Channels.txt")
$domain = [System.DirectoryServices.ActiveDirectory.Domain]::GetCurrentDomain()
    "This computers domain information is:" >> $OutputFileName
$domain >> $OutputFileName
    "This computers secure channel information is:" >> $OutputFileName
gwmi Win32_NTDomain    >> $OutputFileName

```

Listing 2: Detect and Report Secure Channel Problems

```

PowerShell
$Domain = "americas"
function SecureChannelCheck
{
    #Function to give a simple "good" or "bad" result for secure channel health.
    #Accepts a flat domain name--not entire FQDN--as input.
    #To run as a script, not a function, just replace $DomainName with
        $env:userdomain.
    param ($DomainName)
    $v = "select * from win32_ntdomain where domainname = '" + $DomainName + "'"
    $v2 = get-wmiobject -query $v

    if ($v2.Status -eq "OK")
        {Write-Host "The domain secure channel is OK."}
    elseif (($v2 -eq $null) -or ($v2 -ne "OK"))
        {Write-Host "The domain secure channel has problems."}
}

SecureChannelCheck($Domain)

```

Listing 1 obtains secure channel information for the current domain, and basic information about the forest. (Figure 2 shows the results.) To detect any problems, we simply create our test as a PowerShell script (.ps1 file) and add an If statement to the returned status, and we can even supply the domain name, as in Listing 2. For the Microsoft diagnostic case, we also make it into a simple function so that we can reuse it.

Figure 2: Results from Listing 1

Forest:	contoso.com
DomainControllers:	{BYC-NA-DC-50.americas.contoso.com, CGY-NA-DC-50.northamerica.corp.micros NYC-NA-DC-50.americas.contoso.com, CHI-NA-DC-50.americas.contoso.com...}
Children:	{}
DomainMode:	Windows2008R2Domain
Parent:	contoso.com
PdcRoleOwner:	HQ1-NA-DC-02.americas.contoso.com
RidRoleOwner:	HQ1-NA-DC-02.americas.contoso.com
InfrastructureRoleOwner:	
Name:	americas.contoso.com

Detecting secure channel problems in an enterprise environment is the tough part. Fixing them can be much easier. Hopefully, this article helps give you some tools to easily find those problems when they're happening in your environment. ■

Translating Active Directory Object Names Between Formats

How to easily use the NameTranslate object in PowerShell

Active Directory (AD) administrators are well aware that there are multiple names for individual AD objects. For example, users log on to the system using a logon name (username or domain\username) or user principal name (UPN—username@domainname), but there are other names for the account object, such as its distinguished name (DN—e.g., CN = Joe User,OU = Users,DC = fabrikam,DC = com) or canonical name (e.g., fabrikam.com/Users/Joe User). Because AD objects have multiple names, it's helpful to be able to translate between them. To meet this need, Microsoft created the NameTranslate COM object, which quickly translates AD object names between different name formats.

Using the NameTranslate Object in VBScript

The NameTranslate object is fairly straightforward to use in VBScript. Listing 1 shows a bare-bones example of how you would use the object to translate a name in NT4 format (i.e., domain\username) into DN format.

The code in Listing 1 initializes the NameTranslate object to locate a global catalog (GC), sets the NT4 name (NT4Name), and gets the



Bill Stewart

is a scripting guru who works for Indian Health Service in Albuquerque, New Mexico. He's a contributing editor for *Windows IT Pro* and a moderator for Microsoft's Scripting Guys forum. He offers free tools on his website.



Email



Website

Listing 1: Using the NameTranslate Object in VBScript

```

Const ADS_NAME_INITTYPE_GC = 3
Const ADS_NAME_TYPE_NT4 = 3
Const ADS_NAME_TYPE_1779 = 1

Dim NameTranslate
Set NameTranslate = CreateObject("NameTranslate")

NameTranslate.Init ADS_NAME_INITTYPE_GC, ""

Dim NT4Name, DN

NT4Name = "fabrikam\pflynn"

NameTranslate.Set ADS_NAME_TYPE_NT4, NT4Name

DN = NameTranslate.Get(ADS_NAME_TYPE_1779)

' CN=Phineas Flynn,OU=Engineers,DC=fabrikam,DC=com
WScript.Echo DN

```

object's DN. The [NameTranslate documentation](#) describes the object's methods and numeric constants used in this code.

Using the NameTranslate Object in PowerShell

Unfortunately, using the NameTranslate object isn't straightforward in [Windows PowerShell](#) because the object lacks a type library that tells PowerShell what its properties and methods are. The PowerShell session in Figure 1 illustrates the problem. You can create the object, but you can't call its Init method directly like you can in VBScript. (If the NameTranslate COM object had a type library, the code in Figure 1 would work as expected.)

Fortunately, the Microsoft .NET Framework provides an alternative way to call a COM object's methods using the InvokeMember

```

Windows PowerShell
Copyright (C) 2009 Microsoft Corporation. All rights reserved.

PS C:\> $NameTranslate = New-Object -ComObject "NameTranslate"
PS C:\> $NameTranslate.GetType()

IsPublic IsSerial Name                                     BaseType
-----
True     False     __ComObject                                             System.MarshalByR...

PS C:\> $NameTranslate | Get-Member

TypeName: System.__ComObject

Name                MemberType Definition
-----
CreateObjRef        Method      System.Runtime.Remoting.ObjRef CreateOb...
Equals              Method      bool Equals(System.Object obj)
GetHashCode          Method      int GetHashCode()
GetLifetimeService  Method      System.Object GetLifetimeService()
GetType             Method      type GetType()
InitializeLifetimeService Method      System.Object InitializeLifetimeService()
ToString            Method      string ToString()

PS C:\> $NameTranslate.Init(3, "")
Method invocation failed because [System.__ComObject] doesn't contain a method
named 'Init'.
At line:1 char:20
+ $NameTranslate.Init <<<< (3, "")
+ ~~~~~
+ CategoryInfo          : InvalidOperation: (Init:String) [], RuntimeExcep
tion
+ FullyQualifiedErrorId : MethodNotFound

PS C:\>

```

Figure 1

Problem Encountered
When Using the
NameTranslate Object
in PowerShell

method. The InvokeMember method allows PowerShell to call a COM object's methods indirectly, even if the COM object lacks a type library. Listing 2 shows the PowerShell equivalent of the VBScript code in Listing 1.

Listing 2: Using the NameTranslate Object in PowerShell

```

$ADS_NAME_INITTYPE_GC = 3
$ADS_NAME_TYPE_NT4 = 3
$ADS_NAME_TYPE_1779 = 1

$NameTranslate = New-Object -ComObject "NameTranslate"

$NameTranslate.GetType().InvokeMember("Init",
    "InvokeMethod", $NULL, $NameTranslate,

```


Listing 2: *continued*

```

($ADS_NAME_INITTYPE_GC, "") | Out-Null

$NT4Name = "fabrikam\pflynn"

$NameTranslate.GetType().InvokeMember("Set",
    "InvokeMethod", $NULL, $NameTranslate,
    ($ADS_NAME_TYPE_NT4, $NT4Name)) | Out-Null

$NameTranslate.GetType().InvokeMember("Get",
    "InvokeMethod", $NULL, $NameTranslate,
    $ADS_NAME_TYPE_1779)

```

As you can see, the code in Listing 2 is harder to read than its VBScript equivalent because of the call to the .NET `InvokeMember` method. You can make this code somewhat easier to read by using a function that implements the `InvokeMember` method. Listing 3 is the equivalent of Listing 2, except that it creates the `Invoke-Method` function to handle the calls to the `InvokeMember` method.

In the `Invoke-Method` function, you might notice something a bit odd: a seemingly unnecessary *if* statement at the end. If you run the code in Listing 3 without this *if* statement, PowerShell will appear to output a string (the translated name), but in fact, the output will be an object of type `System.RuntimeType`. This *if* statement prevents PowerShell from including an unnecessary null object in its output stream.

Another thing to note about the code in Listing 3 is that when it calls the `Invoke-Method` function, it encloses the function's third parameter in parentheses. The `Invoke-Method` function's third parameter, which corresponds to the `InvokeMember` method's fifth parameter, can be an array if the method it's invoking requires more than one parameter. The parentheses are technically not necessary in the final line of code, but I included them to be consistent.

Listing 3: Using a Function to Implement the InvokeMember Method in PowerShell

```

$ADS_NAME_INITTYPE_GC = 3
$ADS_NAME_TYPE_NT4 = 3
$ADS_NAME_TYPE_1779 = 1

function Invoke-Method([__ComObject] $object,
    [String] $method, $parameters) {
    $output = $object.GetType().InvokeMember($method,
        "InvokeMethod", $NULL, $object, $parameters)
    if ( $output ) { $output }
}

$NameTranslate = New-Object -ComObject "NameTranslate"

Invoke-Method $NameTranslate "Init" ($ADS_NAME_INITTYPE_GC,
    "")

$NT4Name = "fabrikam\pflynn"

Invoke-Method $NameTranslate "Set" ($ADS_NAME_TYPE_NT4,
    $NT4Name)

Invoke-Method $NameTranslate "Get" ($ADS_NAME_TYPE_1779)

```

Making Things Easier

As you've seen, it's possible to use the NameTranslate object in PowerShell, but it's a nuisance because you can't call its methods directly. This nuisance is mitigated somewhat by using a function like Invoke-Method, which makes the code slightly easier to read and understand. However, "slightly easier" wasn't good enough for me, so I encapsulated the power of the NameTranslate object in a script. Translate-ADName.ps1 provides easy, cmdlet-like access to the NameTranslate COM object.

Download[Download the code](#)

Introducing Translate-ADName.ps1

Translate-ADName.ps1, which you can download by clicking the Download button, requires PowerShell 2.0 or later. The script's syntax is as follows:

```
Translate-ADName.ps1 [-OutputType] <String>  
    [-Name] <String[]> [-InputType <String>  
    [-InitType <String>] [-InitName <String>  
    [-ChaseReferrals] [-Credential <PSCredential>]
```

The -OutputType parameter is required and specifies the name format for the script's output. It's the first positional parameter, so its parameter name (-OutputType) isn't required. Table 1 lists the possible values for the -OutputType parameter.

The -Name parameter is also required and specifies the name to translate. You can't use wildcard characters, but you can specify more than one name if you separate them with commas. If a name contains spaces, you must enclose it in single quotes (') or double quotes ("). The -Name parameter is the second positional parameter, so its parameter name (-Name) isn't required. The -Name parameter also supports pipeline input.

The -InputType parameter specifies the format of the names used with the -Name parameter. Table 1 lists the possible values for the -InputType parameter. Note that -InputType supports two additional values compared to -OutputType. The -InputType parameter's default value is *unknown* (i.e., the system will estimate the format).

The -InitType parameter specifies how the script will initialize the NameTranslate object. This parameter can be one of three possible string values:

- *GC*. The script will locate a GC server and use it to translate names. This is the default value.
- *domain*. The script will use the domain name specified by the -InitName parameter to translate names.

- *server*. The script will use the server name specified by the `-InitName` parameter to translate names.

The `-InitName` parameter specifies the name of the domain or server to use to translate names. The `-InitName` parameter is ignored if `-InitType` is `GC`, but it's required if `-InitType` is either *domain* or *server*.

Table 1: Possible Values for the `-OutputType` and `-InputType` Parameters

Value	Description	Example
1779	RFC 1779 name	CN=Phineas Flynn,OU=Engineers,DC=fabrikam,DC=com
DN	Distinguished name (DN); same as 1779	CN=Phineas Flynn,OU=Engineers,DC=fabrikam,DC=com
canonical	Canonical name	fabrikam.com/Engineers/Phineas Flynn
NT4	Domain\username	fabrikam\pflynn
display	Display name	Flynn, Phineas
domainSimple	Simple domain name	pflynn@fabrikam.com
enterpriseSimple	Simple enterprise name	pflynn@fabrikam.com
GUID	Globally unique identifier (GUID)	{95ee9fff-3436-11d1-b2b0-d15ae3ac8436}
UPN	User principal name (UPN)	pflynn@fabrikam.com
canonicalEx	Extended canonical name	fabrikam.com/Engineers Phineas Flynn
SPN	Service Principal Name (SPN)	www/www.fabrikam.com@fabrikam.com
unknown (-InputType only)	Unknown name format (the system will estimate)	
SIDorSIDhistory (-InputType only)	Security Descriptor Definition Language (SDDL) string for the SID or one from the object's SID history	S-1-5-21-123456789A-123456789-123456789-12345

If you specify the `-ChaseReferrals` parameter, the script enables referral chasing. See the TechNet article “[LDAP Referrals](#)” for more information about referrals. The default is to not chase referrals.

The `-Credential` parameter specifies the credentials to use when translating names. This is useful when the account running the script can’t access the domain, such as when you’re logged on using a computer’s local account rather than a domain account. This parameter uses a `PSCredential` object, which securely stores the credential’s password. Even though the password is stored securely in the `PSCredential` object, the `Translate-ADName.ps1` script has to temporarily decode the password in memory before using it, so there’s a remote possibility of password exposure if this portion of memory gets written to disk during a hibernation operation or system crash.

Using Translate-ADName.ps1

Let’s look at a few examples of how you use `Translate-ADName.ps1` to translate AD object names between different name formats. If you want to translate a single canonical name into a DN, you’d use:

```
Translate-ADName -OutputType DN `
  -Name "fabrikam.com/Engineers/Phineas Flynn"
```

Note that you can omit the `-OutputType` and `-Name` parameter names because they’re the script’s first two positional parameters, like this:

```
Translate-ADName DN "fabrikam.com/Engineers/Phineas Flynn"
```

Because you didn’t specify the `-InputType` parameter, the script uses the default value *unknown* (i.e., the `NameTranslate` object will estimate the input name format). Also, since you didn’t specify `-InitType`, the script uses `GC` as the default value.

This command is rather useful when you’ve opened an object in the Microsoft Management Console (MMC) Active Directory Users

and Computers snap-in and you want its DN. When Advanced Features is enabled, the Object tab displays the canonical name of the object. You can copy the canonical name to the clipboard and use `Translate-ADName.ps1` to easily translate it into a DN.

To translate a name using a specific server, you'd use a command like:

```
Translate-ADName canonical FABRIKAM\pflynn `
-InitType server -InitName fabdc1
```

In this case, the script will use the server named `fabdc1` to get the canonical name of the specified user object.

If you have a lot of names to translate, you can put them into a text file (one name per line) and use the `Get-Content` cmdlet to retrieve each name. For example, if you want to use specific credentials to translate a list of DNs into canonical names, you'd use a command like:

```
Get-Content Names.txt | Translate-ADName `
-OutputType canonical -InputType DN `
-Credential (Get-Credential)
```

When you run this command, PowerShell will first prompt you for the credentials it should use because the `Get-Credential` cmdlet runs first. After you enter the credentials, the script will retrieve the contents of the file `Names.txt`, which you're saying contains a list of DNs (*-InputType DN*). The parameter *-OutputType canonical* tells the script to output canonical names.

Suppose that you need to list the names of all the user accounts in your domain and they need to be sorted by container. The AD cmdlet `Get-ADUser` makes it easy to retrieve a list of the DNs of user objects in the current domain:

```
Get-ADUser -Filter 'Name -like "*"' |
Select-Object -ExpandProperty DistinguishedName
```

Unfortunately, using the `NameTranslate` object isn't straightforward in PowerShell because the object lacks a type library.

But what if you want to sort this list by container? You can't, because the CN attribute appears first in a DN. Consider the following DNs:

```
CN=Phineas Flynn,OU=Engineers,DC=fabrikam,DC=com
CN=Perry the Platypus,OU=Secret Agents,DC=fabrikam,DC=com
```

If you sort these two DNs, Perry the Platypus will get sorted before Phineas Flynn. In order to sort the names by container, you need these names in canonical format, which will now sort properly:

```
fabrikam.com/Engineers/Phineas Flynn
fabrikam.com/Secret Agents/Perry the Platypus
```

By using the Translate-ADName.ps1 script to translate the names provided by Get-ADUser cmdlet, then using Sort-Object cmdlet to sort them, you can get the desired result:

```
Get-ADUser -Filter 'Name -like "*"' |
  Select-Object -ExpandProperty DistinguishedName |
  Translate-ADName canonical | Sort-Object
```

Name Translation Made Easy

The NameTranslate object is a really useful tool when you need to translate AD object names between different formats. Unfortunately, this object is awkward to use from PowerShell because it lacks a type library. The Translate-ADName.ps1 script eliminates this awkwardness and makes it simple to take advantage of this useful object in PowerShell. ■

Mailbox Auditing in Exchange Server 2010

How to audit individual mailboxes using PowerShell

Microsoft added the ability to audit mailbox actions in Exchange Server 2010 SP1 as part of its initiative to make Exchange capable of satisfying the compliance requirements of large organizations. Other compliance features in Exchange 2010 include multi-mailbox discovery searches and search mailboxes, administrator auditing, and retention policies, all of which are included in [Exchange Server 2013](#). Mailbox auditing complements search mailboxes in some ways because these mailboxes usually hold a lot of confidential information that has been extracted from user mailboxes. It's good to have the ability to audit access to these mailboxes to control appropriate access and operations. Mailbox auditing also addresses the common requirements of tracking access to other sensitive mailboxes (e.g., those used by executives) and determining who sent a particular message from a shared mailbox.

After I describe how mailbox auditing is implemented in Exchange 2010 and later, I'll show you how to:

- Enable and configure mailbox auditing so that the right data is collected
- Suppress audits for specific mailboxes
- Search mailbox audit data with Windows PowerShell
- Report audit data with Exchange Control Panel (ECP)
- Get auditing data for heavily loaded servers

All the mailbox auditing features and commands described here work with an on-premises deployment of Exchange 2010 and with Exchange



Tony Redmond

is a contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press).



Email



Twitter



Blog

Online in Microsoft Office 365. They also work with Exchange 2013, with the difference being that you use the Exchange Administration Center (EAC) instead of Exchange 2010's ECP.

In Exchange 2010, there isn't a GUI to control mailbox auditing in either the EMC or ECP. The lack of a GUI continues in Exchange 2013.

Understanding How Mailbox Auditing Is Implemented

Mailbox auditing in Exchange 2010 and later has the following characteristics:

- Auditing is configurable on an individual mailbox basis rather than for a complete server.
- You can enable three levels of access: user, delegate, and administrative.
- You can audit up to 11 different actions, such as create and update. (For a complete list of the actions, see “[Mailbox Audit Logging](#).”) You can audit all, some, or just one.
- Audit items are stored in the mailbox being audited rather than in the event log.
- PowerShell cmdlets are available to edit the audit configuration and investigate audit items.

Exchange 2010 and later versions store mailbox audit items in a hidden Audits subfolder in the Recoverable Items folder in the user's mailbox. No client can see the items in the Audits subfolder because it's hidden by Exchange. However, if you really want to, you can use a program such as [MFCMAPI](#) to see what the items look like.

The Recoverable Items folder seems like a regular folder. However, clients have to perform special processing whenever its contents are accessed to ensure that users can't interact with dumpster data that might be required for compliance purposes. Microsoft Outlook and Outlook Web App (OWA) expose a certain amount of information in the Recoverable Items folder through the Recover Deleted Items feature, which lists all the items in the Deletions subfolder in the Recoverable Items folder. Utilities such as MFCMAPI are able to expose more data by opening other subfolders (e.g., Purges, Versions), but

they can't open the Audits subfolder. I think this is logical because you don't want to enable auditing for a mailbox only to discover that an administrator has used MFCMAPI afterward to wipe out the records that reveal some suspicious behavior.

Because mailbox auditing is configured on a per-mailbox basis, it makes sense for Exchange to store the audit items relating to a mailbox in the mailbox itself. Audit items therefore accumulate in the Audits subfolder. They remain there for 90 days (by default), after which they're removed by the Managed Folder Assistant. You can configure the retention period to be anything up to 24,855 days, but I imagine that the default is probably sufficient to perform occasional monitoring of mailbox access. Situations in which mailboxes are placed under litigation hold might require you to enable auditing for those mailboxes as well as to increase the audit retention period.

Audit items aren't large—typically between 1.5KB and 3KB each—so accumulating audit items for 90 days shouldn't create too much of an overhead for the mailbox. The space taken by the audit items isn't charged against a mailbox's quota.

If you're curious, the presence of the Audits subfolder can be seen using the Exchange Management Shell (EMS). For example, the following command scans the mailbox named VP Mailbox to look for any folder named Audits and reports the number of items and their size. A filter suppresses the display of data for any user-created folder that's also called Audits. The command is:

```
Get-MailboxFolderStatistics -Identity "VP Mailbox" |  
  ? {$_.Name -eq "Audits" -and $_.FolderType -eq "Audits"} |  
  Format-Table Identity, ItemsInFolder, FolderSize -AutoSize
```

As Figure 1 shows, there aren't a lot of audit items in this mailbox, so you can conclude that either auditing has only recently been enabled or this isn't a busy mailbox. Background knowledge about how Exchange is used in your environment is invaluable when it comes to

Figure 1
Finding Out the
Number of Audit Items
in a Mailbox's Audits
Subfolder

Identity	ItemsInFolder	FolderSize
VP Mailbox\Audits	74	122.9 KB (125,860 bytes)

interpreting information extracted from the server.

Depending on the events being audited, a busy mailbox that's managed by a delegate might generate hundreds of audit items daily.

Enabling and Configuring Mailbox Auditing

Mailbox auditing isn't enabled by default, so the first thing you need to do is enable auditing for the mailboxes for which you need to collect data. In Exchange 2010, there isn't a GUI to control mailbox auditing in either the Exchange Management Console (EMC) or ECP. (The lack of a GUI continues in Exchange 2013.) Therefore, you need to use the Set-Mailbox cmdlet. For example, this command enables auditing for the mailbox named CEO Mailbox:

```
Set-Mailbox -Identity 'CEO Mailbox' -AuditEnabled $True
```

As soon as auditing is enabled on a mailbox, Exchange starts to generate audit items based on the mailbox's audit configuration. You can see the configuration by running the command:

```
Get-Mailbox -Identity 'CEO Mailbox' | Format-List Audit*
```

Figure 2 shows this command's output. It tells you that auditing is enabled (AuditEnabled property) and Exchange will keep audit items for 90 days (AuditLogAgeLimit property). The output also reveals the

actions that are being audited for the three levels of access that you can manipulate.

(The actions in the output are truncated by PowerShell.) Exchange enables a default set of audit actions, which you can amend.

Figure 2
Reviewing a Mailbox's
Audit Configuration

```
AuditEnabled : True
AuditLogAgeLimit : 90.00:00:00
AuditAdmin : {Update, Move, MoveToDeletedItems, SoftDelete...}
AuditDelegate : {Update, SoftDelete, HardDelete, SendAs...}
AuditOwner : {}
```

The `AuditAdmin` property defines the administrative audit actions. These aren't actions such as an administrator logging on to a mailbox but rather actions such as importing messages into a mailbox from a PST.

`AuditDelegate` is the most interesting of the five properties because it controls auditing for delegates. A delegate is a user who isn't the primary owner of a mailbox but who has some level of access to open the mailbox and access its contents. For example, an administrative assistant might be given the rights to access a manager's calendar and send mail on behalf of the manager. A legal executive might be assigned full access to a discovery search mailbox so that he or she can review the items retrieved by a multi-mailbox search.

As you can see in Figure 2, the `AuditDelegate` property's default audit actions include `SendAs`. It's important to note that Exchange distinguishes between situations in which a delegate uses `SendAs` (send as if the delegate were the user) and `SendOnBehalf` (include an indication that someone else sent the message). These are two very different ways that a delegate can send a message for another user. Although the `SendAs` action is one of the default actions for `AuditDelegate`, the `SendOnBehalf` action isn't. So, if you want to audit both actions, you need to add `SendOnBehalf` to the list of actions. To do so, you can use the `Set-Mailbox` cmdlet to change the audit actions for the `AuditDelegate` property, like this:

```
Set-Mailbox -Identity 'CEO Mailbox' -AuditDelegate `
    "Update, SoftDelete, HardDelete, SendAs, SendOnBehalf, `
    MoveToDeletedItems, Copy"
```

The last property shown in Figure 2 is `AuditOwner`. Unsurprisingly, the default set of audit actions specified for the `AuditOwner` property is empty. This property controls the actions that are audited when the mailbox user accesses items. Auditing every user action will quickly generate a lot of items, so this level of auditing is best reserved for only when it's absolutely necessary.

Suppressing Audits for Specific Mailboxes

You might have situations in which you don't want to generate mailbox audit entries. The most common scenario is when you have a service account that has delegated access to many user mailboxes. For example, Research in Motion's BlackBerry Enterprise Server (BES) for Microsoft Exchange needs to access mailboxes to retrieve new mail to deliver to BlackBerry devices. It uses the same access to replicate operations executed on mobile devices back to mailboxes. The BES service account might generate many audit records daily, which could create a situation in which audit records that require further investigation are hidden by records belonging to a mass of totally innocent and mundane operations. To avoid this, you can run the [Set-MailboxAuditBypassAssociation](#) cmdlet to identify accounts that Exchange should ignore when auditing mailboxes. For example, this command instructs Exchange to ignore any access by a service account named BESService:

```
Set-MailboxAuditBypassAssociation -Identity 'BESService' `
    -AuditBypassEnabled $True
```

Obviously, it's not a good idea to have administrators run this cmdlet so that their own accounts are excluded from mailbox auditing. Exchange doesn't currently provide any method to block administrators from taking such an action, so all you can do is ensure that administrative auditing is enabled for the organization and that the `Set-MailboxAuditBypassAssociation` cmdlet is included in the audit configuration. For example, to include the cmdlet in the set audited by Exchange, you'd run a command similar to this:

```
Set-AdminAuditLogConfig -AdminAuditLogCmdlets `
    "Set-MailboxAuditBypassAssociation, Set-Mailbox, `
    New-Mailbox, *Transport"
```

(Note that I chose an arbitrary set of cmdlets to audit here.)

Searching Mailbox Audit Data with PowerShell

Exchange 2010 and Exchange 2013 provide two EMS cmdlets to search mailbox audit data:

- **Search-MailboxAuditLog** performs an immediate synchronous search across one or more mailboxes and returns information on screen. You'd use the Search-MailboxAuditLog cmdlet if you need to confirm that auditing is enabled and operational for a mailbox or if want to track down a particular audit event. You can also create a report from the data that you extract by piping it to an external file.
- **New-MailboxAuditLogSearch** operates in the background. It extracts mailbox audit data, puts that data into an XML file, and attaches the file to an email message that's delivered to whatever address you choose. You might use this cmdlet to create a report for an external legal investigator. New-MailboxAuditLogSearch will be discussed in detail in the "Getting Auditing Data for Heavily Loaded Servers" section.

To use the Search-MailboxAuditLog cmdlet, you'd use a command like this:

```
Search-MailboxAuditLog -Identity Billing `
-LogonTypes Delegate -ShowDetails -StartDate "1/1/2012" `
-EndDate "1/31/2012" | ft Operation, OperationResult, `
LogonUserDisplayName, ItemSubject, LastAccessed, -AutoSize
```

In this instance, you're searching only one mailbox, so you pass its name as a value to the -Identity parameter. If you want to search several mailboxes, you can replace the -Identity parameter with a comma-separated list of mailbox names passed to the -Mailboxes parameter.

There are a couple of interesting things that you can observe from the sample results in Figure 3. First, an attempt was made to delete an

item, but that attempt failed. The item's subject isn't captured so you don't know exactly what happened here. However, this item serves to illustrate that all operations are logged when auditing is enabled, including those that fail.

Figure 3
Searching Audit Data

Operation	OperationResult	LogonUserDisplayName	ItemSubject	LastAccessed
SoftDelete	Failed	Tony Redmond		28/01/2012 15:44:42
Update	Succeeded	Tony Redmond	Test	28/01/2012 12:26:19
Update	Succeeded	Tony Redmond	Your bill for \$100K is late!	28/01/2012 12:08:19
SoftDelete	Succeeded	Tony Redmond		12/01/2012 16:24:04
SendAs	Succeeded	Tony Redmond	Your bill for \$100K is late!	11/01/2012 20:14:04

Second, the bottom item reports that a message was sent from the mailbox using the SendAs permission. There's no problem with this. However, you can also see that the item was subsequently updated 17 days later. Those who have suspicious minds might wonder why this happened and contemplate whether it was an attempt to cover something up.

To investigate further, you could do another search to focus in on the suspect audit item. This time, you'd want to search for specific operations within a narrow date range by using a variant of the original EMS command:

```
Search-MailboxAuditLog -Identity Billing `
-LogonTypes Delegate -ShowDetails `
-StartDate "1/28/2012 11:59" -EndDate "1/28/2012 12:15" |
? {$_.Operation -eq "Update"} | Format-Table
```

Figure 4
Investigating Why an
Item Was Updated 17
Days After It Was Sent

```
Operation : Update
OperationResult : Succeeded
LogonType : Delegate
FolderPathName : \Sent Items
ClientInfoString : Client=MSExchRPC
ClientIPAddress : 196.43.160.191
ClientProcessName : OUTLOOK.EXE
ClientVersion : 14.0.6109.5000
LogonUserDisplayName : Tony Redmond
ItemSubject : Your bill for $100K is late!
DirtyProperties : TextBody
MailboxResolvedOwnerName : Billing
LastAccessed : 01/28/2012 12:08:19
```

An edited version of the output is shown in Figure 4. You can now see that the item was updated in the Sent Items folder and that the TextBody property (i.e., the email message's body) was updated. You can also see that Outlook (or a program that

calls Outlook's libraries, such as MFCMAPI) was used, along with the IP address of the computer that was used to make the update. This information should be sufficient to have a conversation with the user to clear up exactly what happened and why.

Administrative operations (e.g., deletion of items from a mailbox with the Search-Mailbox cmdlet) can be recognized because LogonType will be set to Admin in the audit items. For example, here's a command that searches the audit log entries for the CEO's mailbox to locate hard delete operations (i.e., those that permanently delete items):

```
Search-MailboxAuditLog -Identity 'CEO Mailbox' `
-ShowDetails | ? {$_.Operation -eq "HardDelete"} |
  Format-Table Operation, LastAccessed, LogonType, `
  LogonUserDisplayName, FolderPathName, `
  ItemSubject - AutoSize
```

As you can see in Figure 5, the results show that an administrator permanently removed one item from the Budgets folder. Unfortunately, through what can only be an oversight in the Exchange code, you're left hanging as to what that item was, because no information is provided about the subject to help you identify it. This omission wasn't addressed in Exchange 2013, but let's hope it'll be addressed in a future update.

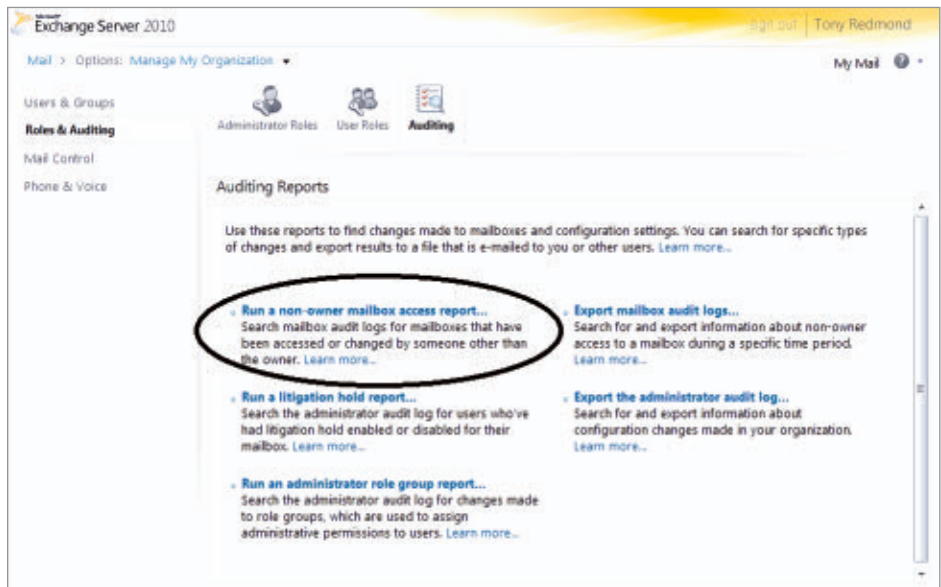
Operation	LastAccessed	LogonType	LogonUserDisplayName	FolderPathName	ItemSubje.
HardDelete	1/31/2012 20:02:14	Admin	AdminUser	\Budgets	

Figure 5
Finding Hard Delete
Operations

Reporting Audit Data with ECP

For those who don't want to use PowerShell to retrieve audit data, Microsoft has provided some out-of-the-box reporting capability for mailbox audit data in the ECP. It's under the Roles and Auditing node of Manage My Organization, as Figure 6 shows. This reporting capability isn't perfect by any means, but it's all mouse-and-click driven.

Figure 6
Using ECP's Out-of-the-Box Reporting Capability for Mailbox Audit Data



By its very nature, a GUI is constrained by the imagination of its designers and the ability of the programmers who write the code. I don't think that the Exchange 2010 developers had an imagination deficit, but the allotted time and resources likely restricted their ability to deliver a fully developed reporting capability. The auditing reports are neither flexible enough nor detailed enough to satisfy true auditing requirements. Unfortunately, nothing changed in Exchange 2013, because the same reporting capability is provided in this version.

One of the available reports, the Non-Owner Mailbox Access Report, lets you search for operations performed in mailboxes by delegates and administrators. ECP displays the screen shown in Figure 7 to obtain the search criteria, such as the date range, the mailboxes to be searched, and the type of access. Interestingly, Office 365 tenants see an access type called *External users*, which refers to access performed by Microsoft data center administrators. The most popular search will likely be for *Administrators and delegated users* because it includes actions such as messages sent by delegates and items deleted in mailboxes. The screen is easy to navigate, and you can experiment with settings to see

what output is produced by ECP. Remember to click the Search button after you change a setting to force ECP to perform a new search.

Search for Mailboxes Accessed by Non-Owners

*Required fields:

Specify a date range and select the mailboxes to search for. Then select to search for non-owner access by anyone or by users inside or outside your organization. [Learn more...](#)

* Start date
2012 January 18

* End date
2012 February 2

Search these mailboxes or leave blank to find all mailboxes accessed by non-owners:
 [Select mailboxes...](#)

Search for access by:
 External users
 All non-owners
 External users
 Administrators and delegated users
 Administrators

[Search](#) [Clear](#)

Search Results

Mailbox	Last accessed:
There are no items to show in this view.	

[Close](#)

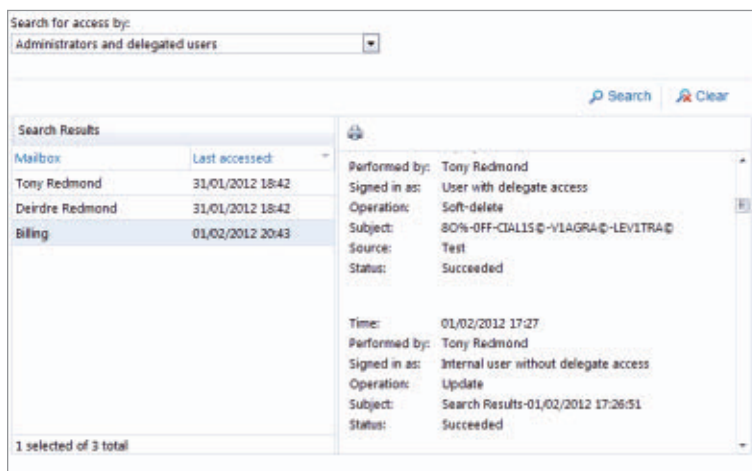
Figure 7

Specifying the Search Criteria for the Non-Owner Mailbox Access Report

Behind the scenes, ECP invokes the Search-MailboxAuditLog cmdlet to retrieve the audit data. It displays the results in the lower right pane, as shown in Figure 8. It'll probably take you a little time to become acquainted with the different kinds of audit data. I recommend that you enable auditing for a mailbox and perform a variety of searches in the mailbox to see what audit information is captured. This will help you interpret audit data more accurately.

For example, take a look at the two audit entries displayed in Figure 8. Both come from the Billing mailbox highlighted in the Mailbox selection pane on the left. The first audit entry reports that a user with delegate access performed a soft delete for an item in the Test folder. This entry is typical of what you'll see when a delegate user accesses a mailbox with Outlook or OWA and deletes a message (i.e., moves it to the Deleted Items folder). The second audit entry is for the same user, but this time

Figure 8
Reviewing the
Retrieved Mailbox
Audit Data



the user has signed in without delegate access to perform an update. You'll see this kind of entry generated by administrative operations run through EMS. In this case, the user ran the `Search-MailboxAuditLog` cmdlet to perform a search and directed the output to a folder that didn't exist. Exchange created the folder and dutifully recorded that action because the folder was in a mailbox subject to auditing.

Getting Auditing Data for Heavily Loaded Servers

It's easy to retrieve auditing data for lightly loaded servers, but it can be quite a different matter for heavily loaded servers, where auditing is enabled on many mailboxes. Exchange provides the `New-MailboxAuditLogSearch` cmdlet for this purpose. (ECP also uses the `New-MailboxAuditLogSearch` cmdlet for its *Export mailbox audit logs* option.) Running `New-MailboxAuditLogSearch` forces Exchange to execute a background search and return the results in the form of an XML-formatted attachment that's emailed to the specified recipients. Take, for example, the following command:

```
New-Mailbox-AuditLogSearch `
    -Name "Check for Delegated Sends" `
    -LogonTypes Delegate `
```



```
-StartDate '1/1/2012' -EndDate '2/1/2012' `
-StatusMailRecipients "Tony.Redmond@contoso.com"
```

In this case, no mailbox names are specified, so Exchange will search all mailboxes that have been enabled for auditing. (For an Office 365 search, you'd need to add *-ExternalAccess \$False* or *ExternalAccess \$True* to indicate whether you want to include audits for data center administrators.) The *-StatusMailRecipients* parameter specifies the email addresses of the recipients of the report in SMTP format. They can be internal recipients or external recipients to accommodate the situation where you might have to generate an audit report as part of a compliance action that's overseen by external legal advisors.

A mailbox assistant executes the command in the background, and the report eventually turns up in the inbox of the specified recipients. The delay is about 30 minutes in an on-premises deployment but much longer for Office 365. A test that I ran took nearly 12 hours to generate a report that arrived at 1:30 A.M., so Microsoft has probably tuned Office 365 to execute this kind of background processing at times of low system demand. Figure 9 shows a sample XML attachment.

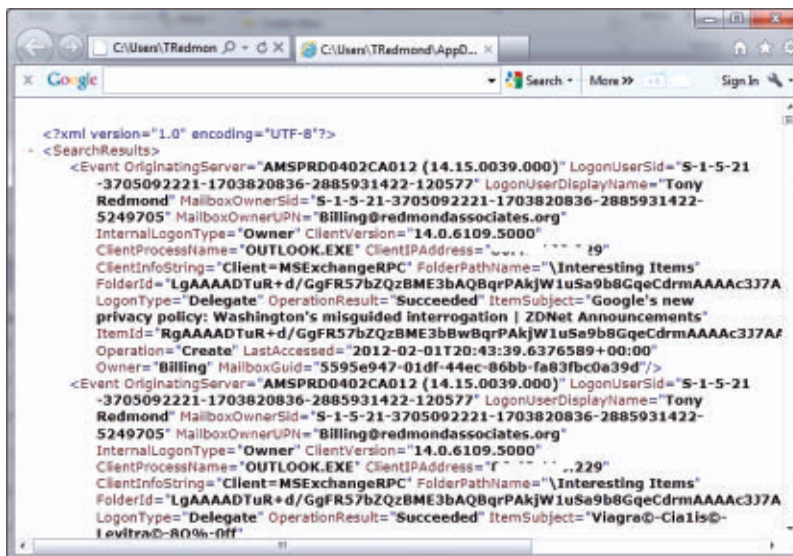


Figure 9
Retrieving Mailbox
Audit Data in XML
Format

Although I have a high regard for the amazing things you can do with XML, I think the decision to provide mailbox audit reports in XML format is flawed. Most recipients won't be fluent in XML, so their first reaction on opening an XML attachment like the one in Figure 9 is likely to be terror rather than pleasure—if they're even able to open the XML attachment. Many clients block this format because of the havoc that can be wreaked by a malicious attachment.

There are better alternatives, but you'll need to use EMS. You can:

- Use the Search-MailboxAuditLog cmdlet to locate the audit data in which you're interested, then build a report using whatever formatting capabilities you desire. There are many different PowerShell examples that show how to generate HTML reports. For instance, Don Jones provides examples in the *TechNet Magazine* article “[Windows PowerShell: HTML Reports in PowerShell](#)” and his free eBook *Creating HTML Reports in PowerShell*.
- Use an already developed solution, such as the excellent solution published by [Andy Grogan on MSEXchange.org](#).

No matter which approach you take, it'll take work to set up the solution. The solution will also require maintenance as Exchange service packs and new versions appear. At a minimum, you'll need to make sure that the code continues to run properly.

A Good Foundation for Mailbox Auditing

Microsoft has built a good foundation for mailbox auditing in Exchange 2010. Audit items are collected as you'd expect, and you have good control over audit configurations. The choice to store audit items in a mailbox is logical, and it's good to see that the items are automatically cleared out after a set period. All that's really limited is the reporting and, to some degree, the search functionality. It would be nice if these aspects of mailbox audit logging are addressed in a future version of Exchange. It's not in Exchange 2013, but perhaps something will come in “Exchange 16.” ■

Windows Server 2012 Installation Options

New configuration levels and on-demand features are a big win

Windows Server 2008 was a huge release, not because of features such as Failover Clustering or Hyper-V, but because of how Windows Server is installed and managed. Changes to these processes in Windows Server 2008 affected every aspect of every installation.

Windows Server 2008 introduced three major technologies that have continued to evolve and become even more important with each new release—especially [Windows Server 2012](#):

- Server Manager
- Windows PowerShell
- Server Core

In this article, I focus on Server Core and some of its complementary technologies in Windows Server 2012.

A Brief History of Server Core

Server Core became an installation choice with Windows Server 2008. Installers could select either a Server Core installation or full installation. After initial selection, this setting couldn't be changed without reinstalling the OS. Both installations provided the same Windows OS with the same kernel, performance capabilities, registry, and authentication. However, Server Core provided a minimal OS environment that included only those components that were required to enable the server to run key infrastructure roles.



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE: Server Infrastructure for Windows Server 2012 and Private Cloud. He's a senior contributing editor for *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).



Email



Twitter



Website

In addition, the functionality of these roles was enhanced by many supported features.

What was missing from Server Core was the .NET Framework (which also meant no PowerShell), Windows Explorer (which meant no Start menu, system tray, or taskbar), Microsoft Management Console (MMC), Server Manager, Internet Explorer (IE), and two Control Panel applets (International and Date/Time). On the plus side, Notepad, the registry editor, and Task Manager were still available. For more details about the issues surrounding the Windows Server 2008 version of Server Core, see the sidebar “[Server Core’s Past](#).”

Fortunately, Windows Server 2008 R2 moved Server Core forward in every way. Not only did Server Manager become remoteable (i.e., it could be used to remotely manage a Server Core installation), but the .NET Framework was broken up, allowing a subset of .NET Framework 2.0, 3.0, or 3.5 to be made available on Server Core. This change enabled support for PowerShell on Server Core and a subset of ASP.NET for Microsoft IIS. In addition, Server Core supported the Active Directory Certificate Services (AD CS) role and File Server Resource Manager (FSRM). In addition, many more management agents began to run on Server Core. (In fact, being able to run on Server Core became a logo requirement for management agents.)

What didn’t change between Windows Server 2008 and Windows Server 2008 R2 was Server Core’s focus: It was still targeted at specific infrastructure roles. And the choice of Server Core or a full installation had to be set at installation time and couldn’t be changed without reinstalling the OS. Server Core was no option for application roles, and many organizations struggled with adopting it because of concerns about initial server configuration, lack of a GUI, and possible problems troubleshooting without a GUI.

Windows Server 2012 Server Core

Windows Server 2012 takes Server Core mainstream, emphasizing its focus as the configuration level of choice, whenever possible, for

Windows Server installations. That's right: We now think in terms of configuration levels for Windows Server. Server Core is still one option, as is a full installation—now called the *Server with a GUI* configuration level—but there are others. Server Core is now the default configuration level. Rather than thinking of Server Core as a lesser OS, think of it as the core server OS, available in addition to a version with a GUI.

There are many improvements in Server Core with Windows Server 2012. Additional supported infrastructure roles and role services include Windows Software Update Services (WSUS), Active Directory Rights Management Services (AD RMS), Routing and Remote Access Services (RRAS), Remote Desktop Connection Broker (RD Connection Broker), Remote Desktop Virtualization Host (RD Virtualization Host), and Remote Desktop Licensing (RD Licensing). Nearly all Windows Server roles are now supported on a Server Core installation.

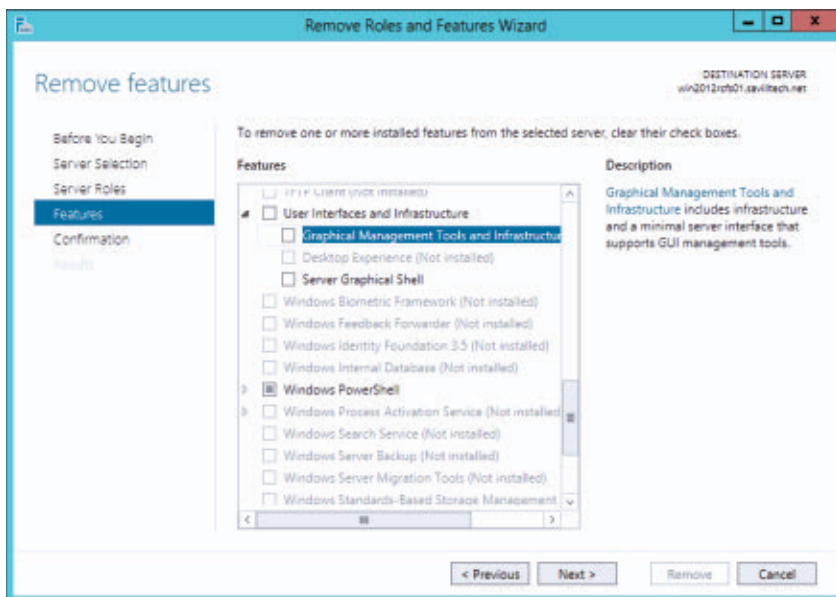
In its position as the preferred configuration level for servers, Server Core is no longer just for infrastructure roles. With Windows Server 2012, Server Core can also be thought of as a platform for running applications. Many applications are already making this switch, including SQL Server 2012, which now runs on a Server Core installation. As I mentioned earlier, one logo criteria for Windows Server 2012 applications is a minimal server interface that allows execution on Server Core. This alone should drive application vendors to add full support for Server Core in future versions of their products.

Windows Server 2012 puts a much greater emphasis on remote management. This is why the actual server OS doesn't need a rich GUI locally. Look at both Microsoft and third-party products: The provided management interfaces are becoming graphically richer and richer, which can consume a lot of resources and therefore ideally should not be running on the server anyway. The design goal of Windows Server 2012 is that it should be managed remotely from a Windows 8 client or a dedicated management Windows Server 2012 installation. This remote-management design goal explains why remote management is enabled by default on a fresh Windows

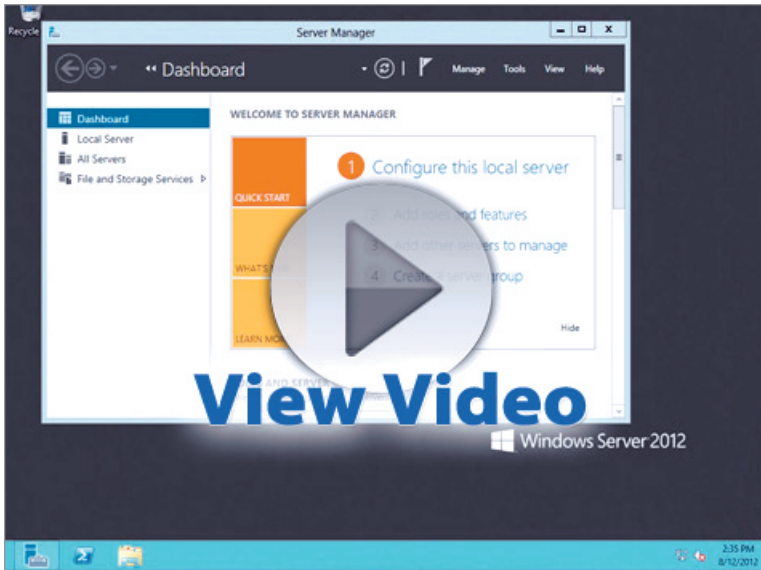
Server 2012 installation, allowing remote management without additional configuration. (Note that *remote management* means that the server can be managed remotely, *not* that Remote Desktop Services is enabled by default.)

The emphasis on Server Core, the ability to use it as an application platform in addition to an infrastructure platform, and the benefits of reduced patching are all great. But remember that many organizations' fundamental objections to Server Core were the initial setup of the OS (using the command line), ongoing management effort, and troubleshooting situations. For these organizations, there is good news in Windows Server 2012. The GUI and management tools can be added or removed from a server installation at any time, requiring only a reboot to complete the switch. The installation choice between Server Core and *Server with a GUI* is no longer set in stone. Figure 1 shows the Remove Roles and Features Wizard as I remove the *Graphical Management Tools and Infrastructure* and the *Server Graphical Shell* features from my *Server with a GUI* installation. After this server reboots, it will be at the Server Core configuration level. None of the

Figure 1
Management Tools
and Graphical Shell



applications or configuration will be lost. You can see more of this process in the accompanying video.



Video

John Savill demonstrates how to switch between a GUI and Server Core installation

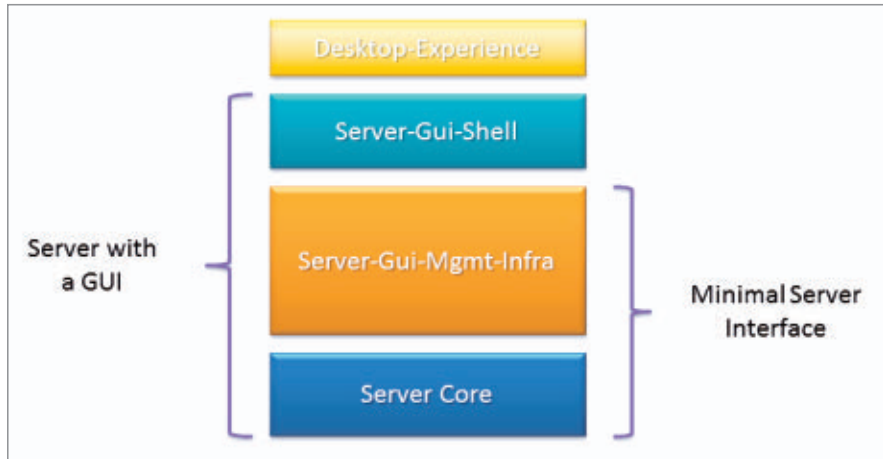
One change in Server Manager in Windows Server 2012 (other than the completely redesigned UI and the ability to manage groups of servers) is that roles and features can be added and removed remotely. Now the switch from Server Core to *Server with a GUI* and vice versa can be made remotely. Even if a server is in Server Core configuration, it can be switched to *Server with a GUI* by using Server Manager, running remotely. This task can also be accomplished from PowerShell and the command line. But before I go into those details, I want to elaborate on configuration levels.

Configuration Levels

I mentioned earlier that there are more configuration levels than just Server Core and *Server with a GUI*. There are actually four configuration levels, although only three are used for most server installations. These configuration levels are shown in Figure 2, along with the components that you can add to Server Core to reach the various levels.

Figure 2

Four Levels of
Configuration for
Windows Server 2012
Installation



The Server Core component forms the base installation or configuration level, providing the key Windows services. On top of this foundation, various components can be added to change the configuration level. For example, adding the Server-Gui-Mgmt-Infra and Server-Gui-Shell components brings a server to the *Server with a GUI* configuration level, providing the full graphical and local management-tool experience.

This separation of the graphical shell from the management infrastructure provides additional options. Although the Server-Gui-Shell component includes Explorer.exe, IE, and the Start Screen (i.e., the graphical shell), it is the Server-Gui-Mgmt-Infra component that contains all the management tools and management infrastructure (i.e., MMC, Server Manager, and other internal components that many applications rely on). A Windows installation that has the management infrastructure but not the graphical shell installed is referred to as a Minimal Server Interface installation.

In Figure 2, you can see that the Server-Gui-Mgmt-Infra component is larger than the Server-Gui-Shell component. The reality is that Server-Gui-Shell contains only around 300MB of binaries, whereas Server-Gui-Mgmt-Infra contains around 4GB of binaries. This difference gives you a good indication of why a pure Server

Core installation should be the desired end goal. Think of the Minimal Server Interface installation (with Server-Gui-Mgmt-Infra) as a full installation without the shell; the Server Core is the true stripped-down OS installation.

What does this new-found flexibility mean for the Server Core configuration level? Think of prior objections to using Server Core, specifically the pain of initial configuration and troubleshooting without a GUI. (Ongoing management with a GUI is no longer an issue because management will be performed remotely via Server Manager and other graphical tools, as previously discussed.) The ability to add and remove the graphical shell and management infrastructure at any time means that administrators performing manual installations and initial configurations can install a server at the *Server with a GUI* level, perform the configuration, then remove the shell and management infrastructure for the server's normal operations. If during these operations a problem is encountered and needs to be troubleshooted locally on the server (something that shouldn't often occur), and if the lack of local graphical shell and management tools hinders progress, then the administrator can add back the shell and management tools, troubleshoot and resolve the issue, then remove the shell and tools again.

I've not yet mentioned the final component (and configuration level), Desktop-Experience, because very few servers will need it. Desktop-Experience, as the name suggests, gives a server the experience that's associated with the Windows 8 client OS. This experience includes elements such as Windows Media Player and its associated codecs, photo management, and Windows Runtime (WinRT). This latter component allows the execution of Windows 8 Modern (aka Metro) applications that use the new WinRT, including access to the Windows Store. Very few servers need or should want this capability; the exceptions are a specific Microsoft Exchange role that needs the media codecs and people who use the server OS as their desktop OS. (And now that Windows 8 has Hyper-V built in, that shouldn't be

necessary.) Remote Desktop Session Host (RD Session Host) also uses Desktop-Experience to give users a richer experience.

Using PowerShell or the Command Line to Switch Configuration Levels

Figure 1 showed the graphical way to add and remove the various features that make up the configuration levels. You can also use the command line and the Deployment Image Servicing and Management (DISM) command-line tool. Or you can use PowerShell and the `Install-WindowsFeature` and `Uninstall-WindowsFeature` cmdlets. DISM uses slightly different feature names than PowerShell, so stick to using PowerShell when possible.

The following code shows the use of DISM to switch from a *Server with a GUI* installation to Server Core:

```
Dism /online /disable-feature /featurename:ServerCore-FullServer
```

The following code shows the use of DISM to move from Server Core to *Server with a GUI*:

```
Dism /online /enable-feature /featurename:ServerCore-FullServer  
/featurename:Server-Gui-Shell /featurename:Server-Gui-Mgmt
```

The following command uses PowerShell to switch from a *Server with a GUI* installation to Server Core. Note that with this method, only `Server-Gui-Mgmt-Infra` needs to be removed; doing so automatically removes `Server-Gui-Shell`, which is dependent on `Server-Gui-Mgmt-Infra`.

```
Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -Restart
```

To use PowerShell to switch from Server Core to *Server with a GUI*, use this command:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell  
-Restart
```

You can simply add `Server-Gui-Mgmt-Infra` to switch from `Server Core` to a Minimal Server Interface installation:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra -Restart
```

Or you can remove `Server-Gui-Shell` from a *Server with a GUI* installation to switch to a Minimal Server Interface installation:

```
Uninstall-WindowsFeature Server-Gui-Shell -Restart
```

All this flexibility in configuration level is great, but which level is right for you? Aim for `Server Core`, and make sure that you have management agents such as backup, monitoring, and malware that run on `Server Core`. Run all the infrastructure roles and server applications that are supported on `Server Core`.

For server applications that don't run on `Server Core`, use a Minimal Server Interface installation (i.e., use the management infrastructure but not the graphical shell). I've seen many software installers that have GUI and IE dependencies, but once installed, the app runs fine on Minimal Server Interface. If something isn't installing, it might be an installer issue, so convert to *Server with a GUI*, install the application, and then convert back to Minimal Server Interface. If an application has dependencies on the graphical shell, then you'll need to run it on the *Server with a GUI* configuration level. However, many applications should eventually develop management tools that can run remotely, with no dependencies on the graphical shell or management infrastructure.

No matter which configuration level you use for a Windows Server 2012 system, manage it remotely. Your Windows Server 2008 R2 and Windows Server 2008 servers can also be managed remotely, using

**Windows Server
2012 takes
Server Core
mainstream.**

the Windows 8 Server Manager, providing that they have Windows Management Framework 3.0 installed.

Watch the Kilobytes

In England, there's a saying: Watch the pennies, and the pounds will take care of themselves. Basically, this means that if you're responsible with your money at the lowest level, the bigger finances will be fine. The same idea applies to disk space: Watch the kilobytes, and the megabytes will take care of themselves. In other words, be responsible at the server level, and your overall disk usage should be fine. This brings me to another change in Windows Server 2012: Features on Demand.

If you've ever managed Windows NT Server 4.0, then you know the "pleasure" of needing to insert the installation media and reinstall service packs every time a new service is enabled. This routine went away with Windows 2000 Server, when Microsoft decided that the management and time savings were worth the extra disk space for the OS installation. The solution introduced with Windows 2000 was to copy the entire OS and all possible services (now called roles and features) to the server's hard disk, storing them in the WinSxS folder (side-by-side assembly). When a service is enabled, the files are simply taken from the WinSxS folder; no installation media is required. When a service pack is applied to a server, all files in the WinSxS folder are patched, so if a service is later enabled on the server, there is no need to reapply the service pack.

However, in a few very specific instances, reducing the disk footprint of an installation might be more important than saving management effort. Imagine having 300 instances of the same virtual machine in a virtual environment, such as an IIS load-balanced farm. In such a case, the gigabytes of additional disk space used for the (unlikely to ever be needed) binaries of each instance add up to a huge amount of additional required storage. Another example is placing Windows Server 2012 on a small storage device, such as on a

motherboard. In this case, minimizing the disk footprint is crucial. Windows Server 2012 allows roles and features to be removed from the WinSxS folder, reducing the disk footprint. This removal is performed by using the PowerShell `Uninstall-WindowsFeature` cmdlet with the `-Remove` parameter

```
Uninstall-WindowsFeature <FeatureName> -Remove
```

or by using DISM with the `/remove` switch

```
Dism /online /disable-feature /featurename:<FeatureName> /remove
```

If an installation has had a role or feature removed from the disk and then tries to add the role or feature, then by default Windows Server downloads the required binaries from Microsoft Windows Update servers on the Internet. Alternatively, you can specify the WinSxS folder of another server, the Windows Server installation Windows Imaging Format (WIM) file from the installation media, or a share with the contents of the extracted installation WIM file. This can be specified as part of the removal command, set as the default in the server registry (via the `HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Servicing\LocalSourcePath` value), or set using Group Policy.

This functionality should not be used as a default part of server installation but should be reserved for very specific use cases, as I already mentioned. Although saving disk space might seem like a good idea, the additional issues you might encounter if roles or features need to be added in the future can outweigh the disk-space savings unless reducing space is a necessity.

A Rewarding Shift

I don't want to underplay the scale of these changes—especially configuration levels: It's big. Many organizations will still manage each server locally by using RDP and running all configurations directly

Server Core's Past

Server Core was introduced in Windows Server 2008 for a very specific reason. Each Windows version introduces new features that offer a richer OS than before. But these changes also make the OS bigger, with a larger footprint, larger potential attack surface, more to patch, and more associated reboots. In addition, the Windows Server and Windows Client OSs share a large amount of code, including components that might not be needed for a server (e.g., media services, Internet Explorer—IE, the graphical shell Explorer.exe). This adds even more code to the server OS but provides a consistent experience for administrators between the client and server. But when a server is a domain controller (DC), a file server, or a simple web server, all the extra features and even the nice GUI are unneeded; they just add to what needs to be patched. Enter Server Core as the preferred platform for infrastructure servers.

After logging in to Server Core in Windows Server 2008, the administrator would see a command prompt box that was used to manage the server. This was a shift in management for many Windows administrators, who were used to a graphical experience. Because Server Core in Windows Server 2008 was the minimal OS needed to run the key infrastructure roles, it omitted many components that were typically found in Windows Server. Some of those components required more frequent patching. Compared with a full installation, Server Core required, on average, half the number of patches and needed to be rebooted far less frequently. In addition, the Server Core installation had a smaller disk footprint and some reduction in CPU and memory. But the real benefit of using Server Core was the reduction in patching and the resulting improvement in server availability.

You might think, then, that Server Core would have been widely adopted. After all, who wouldn't want to halve the number of patches needed and reduce the associated management? But the other two focus areas for Windows Server 2008—namely Server Manager and Windows PowerShell—give some clues as to why very few companies adopted Server Core.

With Windows Server 2008, Microsoft introduced Server Manager. This was the new way to manage servers; it was role-focused, providing an intuitive experience for administration. However, Server Manager couldn't manage a remote machine and thus couldn't be used with Server Core. As a result, administrators were unable to use the new management direction with Server Core and instead needed to use the legacy Microsoft Management Console (MMC) tools and the command line. Also with Windows Server 2008, Microsoft pushed PowerShell as *the* direction for management of all things Microsoft. PowerShell was built on the .NET Framework. Because of the monolithic nature of .NET (i.e., it being one big component that couldn't be broken up) and its reliance on components that weren't available in Server Core, it was impossible to provide .NET Framework for Server Core. Therefore, PowerShell was unavailable on Server Core. In addition, many third-party management agents (and their backup, monitoring, and malware-protection products) didn't work correctly on Server Core because of dependencies on components that weren't found in that installation. ■

on the server. Remotely managing and using Server Core is a huge change, but organizations that make the shift can reap large rewards with simplified patching and management. ■

What's New in Lync Server 2013

The enhancements will make administrators' and users' jobs easier

Microsoft Lync Server 2013 does a commendable job building on the features introduced in Lync Server 2010, making the transition to this new version compelling for IT managers and unified communications (UC) implementers. Like its predecessor, Lync Server 2013 provides support for enterprise IM, presence, and conferencing, with direct integration to the Microsoft Office suite of applications. (If you're unfamiliar with Lync, presence is a feature that displays a user's availability, willingness to communicate, and contact information.)

There are several significant architectural changes that simplify the deployment of Lync Server 2013, with one of the biggest being server role consolidation. I'll discuss the server role changes as well as other changes to Lync Server 2013's topology. I'll also discuss other noteworthy enhancements made to the:

- Role-based access control (RBAC) feature
- Disaster recovery, high availability, and archiving features
- Edge Server services
- IM and presence features
- Conferencing features
- VoIP functionality, which is called Enterprise Voice
- Lync 2013 client

Lync Topology Changes

The Lync topology is the overall layout of the Lync environment, including Front End Servers, Edge Servers, and Mediation Servers. The



Byron O. Spurlock

is an independent consultant and trainer specializing in unified communications and messaging. He writes a blog about Exchange Server and unified communications.



Email



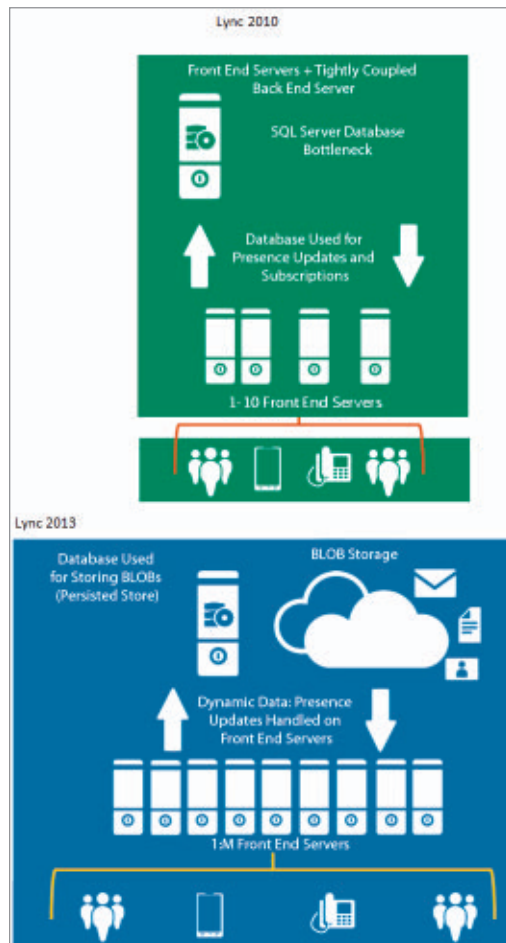
Blog

Lync topology consists of multiple server roles that work together to provide IM, presence, conferencing, and VoIP functionality. The topology changes in Lync Server 2013 bring about many rich enhancements that will affect how administrators design and ultimately deploy their Lync environments and approach upgrades.

One of the most important topology changes deals with the Front End pools. In Lync Server 2013 Enterprise Edition, the architecture of the Front End pools has been changed to provide a more distributed systems architecture, as Figure 1 shows. In the Lync Server 2013 architecture, the back-end Microsoft SQL Server database is no longer the point of reference for the real-time data store in the Lync pool that contains the updated presence information, permissions, and user contacts. Those responsibilities have been transferred to the Lync 2013 Front End Servers. This distribution of data storage improves performance because the Back End Server no longer has to render the up-to-the-second transactions for users regarding presence, contacts, and conferencing information. It also provides scalability within the pool and eliminates the single point of failure (i.e., the single Back End Server).

There are many changes to the server roles in Lync Server 2013. Here are the most important changes:

Figure 1
Changes in the
Architecture



- Lync Server 2013 no longer has a separate Archiving Server role. Archiving is an optional feature available on all Front End Servers.
- Lync Server 2013 no longer has a separate Monitoring Server role. Monitoring is an optional feature available on all Front End Servers.
- The A/V Conferencing Server role is now part of the Front End Server role.
- The Persistent Chat Server is a new server role. (Persistent Chat is the new name for Group Chat. Persistent Chat will be discussed in more detail in the “[Lync 2013 Client Enhancements](#)” section.)
- The Director role is no longer presented as a recommended role but rather an optional role. Organizations that have specific security requirements for allowing external traffic through the perimeter network to the Lync servers inside the network might consider deploying this role.

RBAC Enhancements

In Lync Server 2013, Microsoft enhanced the RBAC feature in two major ways. First, there are two new predefined roles: Response Group Manager and Persistent Chat Manager. Members of the Response Group Manager role can manage specific Response Group queues. Members of the Persistent Chat Manager role can manage specific Persistent Chat rooms.

The second enhancement is the ability to create custom roles. You can create custom roles that have privileges to run a specified set of Windows PowerShell cmdlets or specified PowerShell scripts in the Lync Server Management Shell.

Disaster Recovery, High Availability, and Archiving Changes

As in Lync Server 2010, Lync Server 2013’s main high availability approach is based on server redundancy through pooling. If a server running a certain server role fails, the other servers in the pool running the same role will take the load of that server. This applies to Front End Servers, Edge Servers, Mediation Servers, and Directors.

Lync Server 2013 adds new disaster recovery measures by enabling you to pair Front End pools together. The two pools can be in the same geographic location or different geographic locations. If one of the Front End paired pools goes down, an administrator can manually fail over the users from the primary pool to the backup pool to provide continuation of service.

Lync Server 2013 also adds Back End Server high availability. This is an optional topology in which you deploy two Back End Servers for a Front End pool and set up synchronous SQL Server database mirroring for all the Lync databases running on the Back End Servers. You can also deploy a witness for the mirror.

Enabling the Archiving feature in a Lync 2013 pool provides new capabilities that improve deployment and operation efficiency, including:

- [Microsoft Exchange Server 2013](#) integration. The deprecation of the Microsoft Messaging Queuing Service and the usage of the Lync Storage service in Lync Server 2013 mean that you can use a unified archiving storage architecture if your environment is running Exchange 2013 and Lync Server 2013. When you enable the Archiving feature, you can integrate data storage for Archiving with your existing Exchange 2013 storage for all users who are homed on Exchange 2013 and have their mailboxes put on In-Place Hold. With this new archiving environment, you don't need to deploy separate SQL Server databases to archive Lync data. In addition, you can search and retrieve data from a single database instead of multiple databases when searching for user information for compliance situations.
- SQL Server database mirroring. When you deploy the Archiving feature, you can enable SQL Server database mirroring for your archiving database. This gives you more flexibility if you need to provide high availability across data centers.
- Archiving of whiteboards and polls. Archived conference content now includes whiteboards and polls that are shared during the meeting.

Changes to the Edge Server Services

Lync Server 2013 introduces changes to better integrate and extend the existing Edge Server services that are available to organizations. The following is a high-level overview of the changes that can affect the planning and deployment of Edge Server services.

Support for IPv6 addressing. Lync Server 2013 supports IPv6 addressing for all Edge Server services.

Extensible Messaging and Presence Protocol (XMPP) proxy and gateway. Lync Server 2013 introduces a fully integrated XMPP proxy (which you deploy on Edge Servers) and XMPP gateway (which you deploy on Front End Servers). You can deploy XMPP federation as an optional component. Once deployed and configured in the organization, you can create and configure policies that support XMPP federated domains. Optionally, you can configure relationships with each XMPP federated partner to allow users to place contacts from specific organizations in their contacts list. Doing this can strengthen security.

Mobility services for mobile clients. Mobility services enable supported Lync mobile clients to perform such activities as sending and receiving IM messages, viewing contacts, and viewing presence information right out of the box. In addition, mobile devices support some Enterprise Voice features, such as clicking once to join a conference, single number reach, voicemail, missed call notification, and making and receiving calls on a mobile phone using a work phone number instead of the mobile phone number (i.e., the Call via Work option).

IM and Presence Enhancements

Lync Server 2013 has new IM and presence features. If your organization runs Exchange 2013, users can take advantage of a unified contact store. Users can manage their contacts in Outlook 2013, Outlook Web App (OWA), and the Lync client. In addition, the new XMPP integration feature lets Lync users exchange IM messages and presence information with Google Talk users and users of other public IM providers that utilize XMPP.

In Lync Server 2013 Enterprise Edition, the architecture of the Front End pools has been changed to provide a more distributed systems architecture.

Conferencing Enhancements

Lync Server 2013 introduces many new and updated features that enhance conferencing. The updated Join launcher now validates each meeting before launching a client. It supports opening a meeting in the following clients:

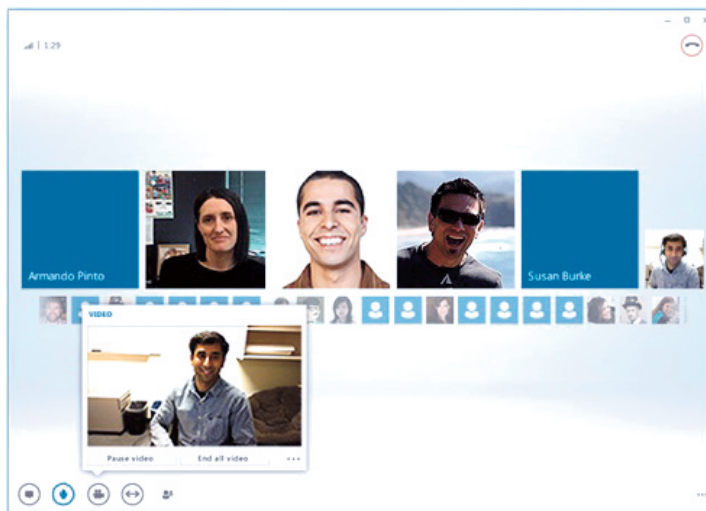
- [Windows 8](#)
- Microsoft Internet Explorer (IE) 10
- Windows Phone 7
- Google Android devices
- Apple iOS devices

Lync Server 2013 supports HD video. Users can experience resolutions up to HD 1080P in two-party calls and multiparty conferences. The H.264 video codec is the default for encoding video on Lync 2013 clients. It supports a greater range of resolutions and frame rates, and improves video scalability. It also allows a more adaptive approach for the Lync infrastructure to render and mix HD resolution to clients.

After users connect to a conference, they can see everyone in the Gallery View. If a video isn't available for a participant, the person's Lync profile picture will appear, as Figure 2 shows. To see the

Figure 2

Gallery View After a User Connects to a Conference



participants' names, users can hover the mouse over the View Participants button or click Show Participant List.

In video conferences with two to four people, users will see the videos of all the participants in the Gallery View (assuming everyone has a video feed). If the conference has more than five participants, videos of only the most active participants appear in the top row and photos appear for the other participants. Presenters can use the video spotlight feature to select one person's video feed so that every participant in the meeting sees only that participant.

Many other enhancements have been made to improve the users' experience during conferencing, including:

- Video is enhanced with face detection and smart framing, so that a participant's video moves to help keep them centered in the frame.
- Easy-to-use audio controls in the meeting room enable users to control audio options, such as mute, unmute, and change device.
- With split audio and video streams, participants can add their video feed to a conference but dial in to hear the meeting if they just want to participate through audio only or want to leverage their Public Switched Telephone Network (PSTN) device.
- When sharing programs, users can select multiple programs to share if they need to work with more than one program.
- Users can switch between content types using the *Share content and lead meeting activities* option. Users can also use the Meeting Content menu to choose which content they want to share.
- Users can merge another open conversation into the meeting by using the Merge This Call Into option on the More Options menu.
- Meeting recordings are automatically saved in a format (e.g., MP4) that plays in Windows Media Player (WMP). Users can easily share the file with anyone or use the Publish feature in Recording Manager to post the recording on a shared location.

Lync Server 2013 uses Office Web Apps and Office Web Apps Server to handle PowerPoint presentations shared during a conference. In Lync

Server 2010, presentations were shared remotely with the Lync Web App. Using Office Web Apps Server provides higher-resolution displays, better support for PowerPoint capabilities, and access to more types of mobile devices. It also gives users (with the appropriate privileges) the ability to scroll through a PowerPoint presentation independent of the presentation itself. For more information about Office Web Apps Server, see [“Lync Server 2013: Introduction to Office Web Apps Server.”](#)

Any document that’s shared during a conference is archived in Exchange 2013 data storage if Exchange integration is enabled with the Archiving feature. This includes PowerPoint presentations, attachments, whiteboards, and polls.

Changes to Enterprise Voice

To enhance Enterprise Voice, Microsoft added and enhanced several routing features. Lync Server 2013 supports multiple trunks between Mediation Servers and gateways. A trunk is a logical association between a port number and Mediation Server with a port number and gateway. This means a Mediation Server can have multiple trunks to different gateways, and a gateway can have multiple trunks to different Mediation Servers. Intertrunk routing makes it possible for Lync Server 2013 to interconnect an IP PBX system to a PSTN gateway or interconnect multiple IP PBX systems. Lync Server 2013 serves as the glue (i.e., the interconnection) between different telephony systems.

Figure 3 illustrates the evolution of Mediation Server integration with third-party PSTN gateways. Microsoft Office Communications Server

Figure 3
Evolution of Mediation
Server Integration
with Third-Party PSTN
Gateways



(OCS) 2007, which is the predecessor to Lync Server 2010, provides a one-to-one (1:1) relationship between a single Mediation Server and a single gateway. Lync Server 2010 provides a one-to-many (1:M) relationship between a single Mediation Server and one or more gateways. Lync Server 2013 provides a many-to-many (M:N) relationship, where multiple Mediation Servers can connect to many gateways. Having multiple Mediation Servers communicate with multiple gateways increases the resiliency for inbound and outbound call routing. Mediation Server routing also makes reliability and disaster recovery planning easier.

Lync Server 2013 also introduces new enhancements to voice routing, such as:

- Enhanced call authorization for call forwarding and simultaneous ringing
- Manager/delegate simultaneous ringing
- Voicemail escape (provides a method for managing voicemails when users are set up for simultaneous ringing on multiple phones)
- Caller ID presentation (provides the administrator the ability to modify the format of the calling party's phone number)
- Conference dial-out for users not enabled for Enterprise Voice

Lync 2013 Client Enhancements

Both administrators and users will appreciate the enhancements in the Lync 2013 client. The most noteworthy changes are the client's integration with Office 2013 setup, new administrative templates, virtual desktop infrastructure (VDI) support, Persistent Chat, and a redesigned conversation window. There are also updates to the Lync Web App, the web-based conferencing client for participants outside of an organization.

Integration with Office 2013 setup. The Lync 2013 client and the Online Meeting Add-in for Lync 2013 (which supports meeting management from within the Outlook messaging and collaboration client) are both included with the Office 2013 Setup program. With

the previous versions of Lync, you had to use the Windows Installer properties to customize and control the Office installation.

New administrative templates. Microsoft recommends that you use the Lync Server Management Shell, which utilizes PowerShell, to enforce policies on Lync 2013 clients. However, for certain situations in which the client settings need to take effect before the user actually logs on to Lync, there are a handful of Lync 2013 Group Policy settings that can be applied. The method for deploying Lync Group Policy settings has changed. In the previous versions of Lync, you need to use the Communicator.adm file to define Group Policy settings. In Lync 2013, you can use the Lync .admx and .adml administrative templates, which are part of the Office Group Policy Administrative Templates.

VDI support. The Lync 2013 client now supports audio and video in a VDI environment. After users connect a video or audio device to their computers, they can connect to the virtual machine (VM), log on to the Lync 2013 client that's running on the VM, and participate in real-time audio and video communication.

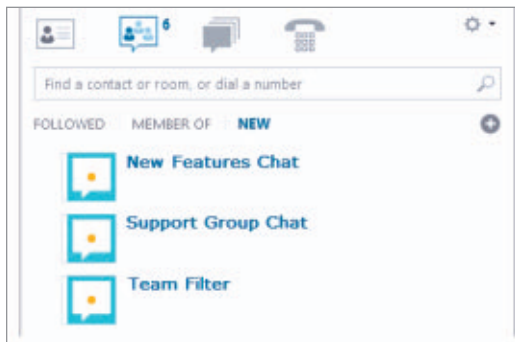
Persistent Chat. The Lync 2013 client integrates the features previously provided by Lync 2010 Group Chat. In other words, a separate Group Chat client is no longer required. In addition, Microsoft changed the name from Group Chat to Persistent Chat. With this single client, users can jump right into a chat room using a UI like that shown in Figure 4. This single client will also make your job easier because you

don't have to deploy and manage multiple clients.

Lync Server 2013 has simplified the administration of Persistent Chat by including an administrative UI that's integrated with the Lync Server Control Panel. Also, the Persistent Chat Server

Figure 4

Persistent Chat
Functionality in the
Lync 2013 Client



includes a collection of PowerShell cmdlets to create and manage Persistent Chat Server categories, rooms (including deleting rooms), and add-ins. You need to be a member of the new CsPersistentChatAdministrator role in order to create and manage chat rooms using the PowerShell cmdlets or Lync Server Control Panel. For more information about Persistent Chat, see “[Lync Server 2013 Persistent Chat](#).”

Redesigned conversation window. The conversation window in the Lync 2013 client has been redesigned to provide quicker access to features. Tabbed conversations are now built into the Lync 2013 client, as shown in Figure 5. With the new tabbed conversations feature, users can keep all their calls, IM messages, and chat rooms in one conversation window. The tabs along the left side of the conversation window let users navigate easily among all active conversations. Other changes include the following:

- With the click of the pop-out button, users can move an individual conversation into a separate window, which they can resize.
- The Lync 2013 client remembers a user’s conversation state, even when the user logs out and logs back on to Lync.
- Users can quickly add IM, video, program sharing, desktop sharing, or web conferencing tools to any conversation by clicking buttons in the conversation window.
- In a meeting where video or content is being shared, users can click the undock button to move the shared video or content into a separate window, which they can resize.

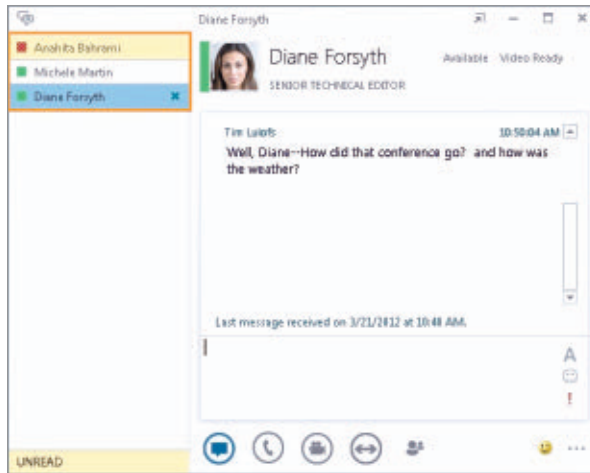


Figure 5
New Tabbed
Conversations Feature
in the Lync 2013 Client

Lync Web App updates. Lync Web App now allows users of Windows or Apple Macintosh PCs to join Lync meetings from within their browser and experience the same features as they would on a regular on-premises Lync client. No local client installation is required to use Lync Web App. In addition, there's no longer a requirement to install Silverlight on the desktop of the user who is participating through the browser. However, some audio, video, and desktop sharing features require the installation of Silverlight. Note that because of the enhancements to Lync Web App, an updated version of Attendee isn't available for Lync Server 2013.

A Top Contender

For users, the changes to Lync's functionality (e.g., IM, conferencing, Enterprise Voice) and client will make Lync 2013 easier and more intuitive to use compared with previous editions. For IT administrators, the changes to Lync Server 2013 (e.g., consolidated roles, RBAC changes, enhanced Enterprise Voice infrastructure) and client will make designing, deploying, and managing the UC environment easier.

The fact that Microsoft hasn't changed some of the basics and is keeping the editions the same is great. You'll still have the option of choosing the Standard or Enterprise Edition, depending on which deployment model best suits your environment. I'm intentionally staying away from discussing sizing numbers until the product becomes a little more solidified, but suffice it to say that the Standard and Enterprise Editions can reportedly host a considerably higher number of users compared with Lync Server 2010.

Being a consultant who does Lync designs and implementations, I can personally say I'm delighted to see the consolidation of the Front End Server roles, the enhanced integration with Exchange 2013, and the evolution of Persistent Chat. The enhancements in Lync Server 2013 are clearly making it a top contender in the UC market. ■

FAQ

Answers to Your Questions

Q: What is the Web Service URL for System Center 2012 Orchestrator?

A: If you deployed System Center 2012 Orchestrator in your environment, other applications will need to know what the Web Service URL is for Orchestrator. The default Web Service URL for System Center 2012 Orchestrator, once deployed, is as follows:

`http://<server name>:81/Orchestrator2012/Orchestrator.svc/`

For example, this is what the Web Service URL looks like in my environment for Orchestrator: `http://savgdalois12.savilltech.net:81/Orchestrator2012/Orchestrator.svc/`.

—John Savill



John Savill



Jan De Clercq

Q: Can Kerberos work across separate Active Directory forests?

A: Yes. If a forest root trust is created between the separate Active Directory (AD) forests, then Kerberos authentication is possible between any domain in any forest because of the transitive nature of the forest root trust. It's very important that services such as DNS are also correctly configured for cross-forest authentication to correctly function. The forest level of both forests must be at least Windows Server 2003. See the Microsoft article "[Creating Forest Trusts](#)" for some key details.

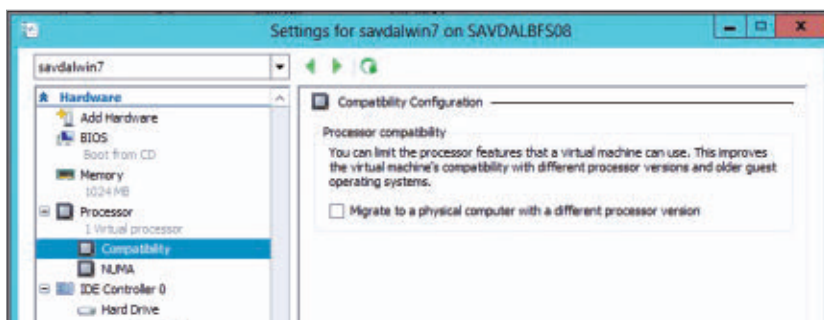
—John Savill

Q: Where is the virtual machine processor compatibility setting in Windows Server 2012 Hyper-V?

A: Windows Server 2012 has the option to enable processor compatibility, which allows a virtual machine (VM) to be moved between Hyper-V hosts that have different versions of the same manufacturer's processor. To enable processor compatibility, perform the following:

1. Start Hyper-V Manager.
2. Open the VM's settings.
3. Expand the Processor node under the Hardware section and select Compatibility.
4. Enable the *Migrate to a physical computer with a different processor version* option (which you can see in Figure 1), then click OK.

Figure 1
Enabling Virtual
Machine Processor
Compatibility



—John Savill

Q: While developing a Windows 8 app using Visual Studio 2012, I got the error “Unable to activate Windows Store app”—what should I do?

A: I’ve seen this happen, and multiple readers have emailed me about this issue. It occurs when you move between machines and try to recompile your application: The application

won't launch and throws the error, which Figure 2 shows.

The easiest solution is to delete the Debug folder from your project's bin folder. Doing so will resolve the problem.




Figure 2
Windows Store App
Error Message

—John Savill

Q: Can I manage Client Hyper-V in Windows 8 from System Center 2012 Virtual Machine Manager SP1?

A: No. Although the Hyper-V feature in [Windows 8](#) is great for many scenarios, it isn't designed to run enterprise services and isn't a supported platform for management by System Center 2012 Virtual Machine Manager.

—John Savill

Q: How can I easily retrieve BitLocker recovery passwords from Active Directory?

A: The Windows BitLocker Drive Encryption Recovery Password Viewer provides an easy solution for retrieving and viewing BitLocker recovery passwords that were backed up to Active Directory (AD). It's an optional feature that's included with Remote Server Administration Tools (RSAT), which you can install by using the *Add Roles and Features* option in the Windows Server 2008 R2 and Windows Server 2012 Server Manager or from the *Programs and Features* option in Control Panel.

The tool provides extensions to the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in and the Active Directory Administrative Center. When the extensions are installed,

you can use the new BitLocker Recovery tab in an AD computer object's Properties dialog box to view the BitLocker recovery passwords that are linked to that computer's BitLocker-protected volumes.

Additionally, you can right-click the domain container in Active Directory Users and Computers and search for a specific BitLocker recovery password across the domain. To do so, in the Find BitLocker recovery password dialog box, which Figure 3 shows, type the first eight characters of the recovery password in the Password ID box, then click Search. Windows displays the first eight characters of the recovery password after the user or Help desk operator reboots a client machine in recovery mode.

Figure 3
Find BitLocker
Recovery Password
Dialog Box



To view the recovery passwords in AD, you must be a domain administrator or you must have been delegated permissions by a domain administrator. The following conditions must also be met: The domain must be configured to store BitLocker recovery information (see “Storing BitLocker recovery information in AD DS” in the Microsoft article [“Backing Up BitLocker and TPM Recovery Information to AD DS”](#) for more information about how to do this), and the BitLocker-protected computers obviously need to be joined to the domain. ■

—Jan De Clercq

Product News for IT Pros

SolarWinds Adds Remediation

SolarWinds announced the latest release of SolarWinds Server & Application Monitor (SAM), the cornerstone of SolarWinds' suite of comprehensive and affordable systems and application management products. Designed to deliver agentless performance and availability monitoring, alerting, and reporting for more than 150 applications and server types, the latest version of SolarWinds SAM will enable IT pros to resolve server and application issues directly from the product's web console. Remediation capabilities include starting and stopping services, killing rogue processes, and rebooting services. Other new and updated features in SolarWinds SAM include new server hardware health monitoring support for popular blade chassis, new native Windows Volume Mount Point monitoring, new Windows Server network performance monitoring for bandwidth usage and errors, and extended integration with System Center 2012 through an updated SolarWinds Management Pack for System Center Operations Manager. For more information about SolarWinds SAM, visit the [SolarWinds website](#).



BiTMICRO Adds PowerGuard Data Protection

BiTMICRO announced that its patented PowerGuard technology would be made available for the company's latest maxIO SSD product line. Although the technology has been available for the company's military SSD products for many years, rapidly increasing enterprise customer demand for device-level power backup has influenced the decision to provide this technology to maxIO drives as well. Recent studies point to a growing concern from enterprise SSD users that data can be lost as a result of power loss or fluctuation. Each PowerGuard-equipped



maxIO SSD device maintains its own internal power backup system. To learn more about PowerGuard technology, visit the [BiTMICRO website](#).

Kodak

KODAK App Simplifies Scanning Documents into SharePoint 2013

KODAK Info Activate Solution Limited Edition delivers the basic essentials from the full version of Info Activate Solution to users who need an easier way to scan documents into Microsoft SharePoint 2013. Available as an app sold and licensed through the Microsoft Office Store, Info Activate Solution Limited Edition is fully integrated with SharePoint 2013 to provide the most efficient method of scanning information from paper documents directly into SharePoint libraries. The Limited Edition is also an excellent way for customers to experience an introduction to the functionality of the full version of Info Activate Solution, offering easy setup and requiring virtually no training. To initiate a specific business process, users click on a large graphic tile that automatically scans and sends a document to its assigned document library in SharePoint 2013. For more information, visit the [KODAK website](#).

Zerto

Zerto Debuts Zerto Virtual Replication 3.0

Zerto introduced Zerto Virtual Replication (ZVR) 3.0, which brings simple, automated disaster recovery to all virtualized workloads at the virtual machine (VM) level. Whether those workloads are mission-critical or tier-two applications, local or remote branch office scenarios, one VM or thousands, based in virtualized data centers or hosted in clouds, ZVR offers 100 percent assurance that all workloads are protected and can be consistently recovered. New enterprise features include robust disaster recovery capabilities for branch offices, a parallel recovery feature to quickly recover large applications, and a new test-before-commit function. The solution also ships with the Zerto Self-Service Portal (ZSSP) and the Zerto Cloud Manager (ZCM). The ZSSP integrates quickly and seamlessly into a service provider's existing

portal, accelerating time to market for disaster recovery services. The ZCM provides a single view of a customer's data, even if the customer is leveraging physical resources across several locations. For more information about ZVR, visit the [Zerto website](#).

ManageEngine Launches Integrated IAM Solution for Windows Shops



ManageEngine launched AD360, the integrated solution for Identity and Access Management (IAM) in organizations running on a Windows-based infrastructure. From everyday user management to keeping user activity compliant with IT regulatory laws, AD360 is engineered to holistically address all aspects of IAM for SMBs and enterprise organizations that use Active Directory (AD). “Without the right tools in hand, accomplishing the IAM goals in a Windows organization becomes a wild goose chase for the IT administrator,” said Manikandan Thangaraj, director of product management at ManageEngine. “AD360 puts the right tools in a single product, creating a unified console that offers an end-to-end solution to [a company's] Active Directory IAM needs.” AD360 unifies ManageEngine's solutions into a modular yet integrated identity and management solution. Users can combine any or all of the AD360 modules to create a single console, single sign-on (SSO), and single product configuration. For more information, visit the [ManageEngine website](#).

Open-E Achieves 600 Percent Improvement Over Active-Passive Failover



Open-E announced that the recently introduced iSCSI (SAN) Active-Active Failover Cluster Feature Pack addition to its Open-E Data Storage Software (DSS) V7 has achieved up to 600 percent performance improvement over existing active-passive failover implementations. The Open-E Active-Active Failover Cluster feature for iSCSI implementations targets applications in high availability, cloud storage, virtualization, and business continuity environments requiring the utmost in

performance and data security. The 600 percent performance improvement using an active-active, load-balanced cluster topology is due to a refinement in the software's algorithm, resulting in a significant decrease in system latency and 100 percent utilization of a system's storage resources. Balancing the workload not only prevents storage overload so that the quality of the service doesn't deteriorate, but it also allows the system to achieve maximum performance, resulting in the cluster's ability to process a high number of read and write operations even faster. For more information, visit the [Open-E website](#).



KineticD Celebrates 10 Years

KineticD announced that 2013 marks the cloud-backup company's 10-year anniversary. Over the past year, the company has demonstrated strong growth, with a 40 percent increase in its reseller channel and a 20 percent uptick in business users worldwide, prompting the opening of a new US data center. "As we enter our tenth year in business, we have seen many changes in, and the rapid development of, the cloud backup and data recovery market," said Jamie Brenzel, CEO of KineticD. "Since its inception, KineticD has been renowned for providing SMBs with the same level of security and protection that is available to large enterprises." The KineticCloud Backup solution is built to support petabyte scalability. More than 60,000 users back up more than 1.2 petabytes of data to KineticD's SSAE16 certified data center. To learn more, visit the [KineticD website](#). ■

7 Groovy USB Gadgets

We searched the Interwebs for our favorite USB gadgets and came up with the wild creations that you see splattered across this page. Click on each image to learn more information and to buy one for yourself!



**Jason
Bovberg**

 Email

 Twitter

 Website



Pure Evil USB 2.0



Star Trek Webcam



USB Pet Rock



i.Saw



USB Squirming Tentacles



8GB Steampunk
Thumb Drive




R2-D2
USB Hub

Submit



Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.

Mobile App
Now Available!



Download your FREE mobile app.

iTunes | Android | Kindle

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowssitpro.com

Support
Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowssitpro.com

News
Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS
Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS
Windows IT Pro VIP
Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowssitpro.com/vip-premium-membership

SQL Server Pro
Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro
Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro
Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

1&1 Internet 1
Windows IT Pro 2, 86

Vendor Directory

Apple 9, 70, 76
BitMICRO 81, 82
BoysStuff.co.uk 85
Geek.com 85
Google 7, 69, 70

i.Saw 85
KineticD 84
Kodak 82
ManageEngine 83
Open-E 83, 84
Research in Motion 44
SolarWinds 81
ThinkGeek 85
Zerto 82, 83

